

INPUT/OUTPUT DEVICES AND INTERACTION TECHNIQUES

Ken Hinckley, Microsoft Research

Robert J. K. Jacob, Tufts University

Colin Ware, University of New Hampshire

INTRODUCTION

The computing literature often draws a sharp distinction between input and output; computer scientists are used to regarding a screen as a passive output device and a mouse as a pure input device. However, nearly all examples of human-computer interaction require *both* input and output to do anything useful. For example, what good would a mouse be without the corresponding feedback embodied by the cursor on the screen, as well as the sound and feel of the buttons when they are clicked? The distinction between output devices and input devices becomes even more blurred in the real world. A sheet of paper can be used to both record ideas (input) and display them (output). Clay reacts to the sculptor's fingers yet also provides feedback through the curvature and texture of its surface. Indeed, the complete and seamless integration of input and output is becoming a common research theme in advanced computer interfaces such as ubiquitous computing (Weiser, 1991) and tangible interaction (Ishii & Ullmer, 1997).

Input and output bridge the chasm between a computer's inner world of bits, and the real world perceptible to the human senses. *Input* to computers consists of sensed information about the physical environment. Familiar examples include the mouse, which senses movement across a surface, and the keyboard, which detects a contact closure when the user presses a key. However, any sensed information about physical properties of people, places, or things can serve as input to computer systems. *Output* from computers can comprise any emission or modification to the physical environment, such as a display (including the cathode ray tube (CRT), flat-panel displays, or even light emitting diodes), speakers, or tactile and force feedback devices (sometimes referred to as *haptic displays*). An **interaction technique** is the fusion of input and output, consisting of all hardware and software elements, that provides a way for the user to accomplish a low-level task. For example, in the traditional graphical user interface, users can scroll through a document by clicking or dragging the mouse (input) within a scroll bar displayed on the screen (output).

The fundamental task of human-computer interaction is to shuttle information between the brain of the user and the silicon world of the computer. Progress in this area attempts to increase the useful bandwidth across that interface by seeking faster, more natural, and more convenient means for users to transmit information to computers, as well as efficient, salient, and pleasant mechanisms to provide feedback to the user. On the user's side of the communication channel, interaction is constrained by the nature of human attention, cognition, and perceptual-motor skills and abilities; on the computer side, it is constrained only by the technologies and methods

that we can invent. Research in input and output centers around the two ends of this channel: the devices and techniques computers can use for communicating with people, and the perceptual abilities, processes, and organs people can use for communicating with computers. It then attempts to find the common ground through which the two can be related by studying new modes of communication that could be used for human-computer interaction (HCI) and developing devices and techniques to use such modes. Basic research seeks theories and principles that inform us of the parameters of human cognitive and perceptual facilities, as well as models that can predict or interpret user performance in computing tasks. Advances can be driven by the need for new modalities to support the unique requirements of specific application domains, by technological breakthroughs that HCI researchers attempt to apply to improving or extending the capabilities of interfaces, or by theoretical insights suggested by studies of human abilities and behaviors, or even problems uncovered during careful analyses of existing interfaces. These approaches complement one another, and all have their value and contributions to the field, but the best research seems to have elements of all of these.

Interaction Tasks, Techniques, and Devices

A designer looks at the **interaction tasks** necessary for a particular application (Foley, Wallace & Chan, 1984). Interaction tasks are low-level primitive inputs required from the user, such as entering a text string or choosing a command. For each such task, the designer chooses an appropriate interaction technique. In selecting an interaction device and technique for each task in a human-computer interface, simply making an optimal choice for each task individually may

lead to a poor overall design, with too many different or inconsistent types of devices or dialogues. Therefore, it is often desirable to compromise on the individual choices to reach a better overall design.

There may be several different ways of accomplishing the same task. For example, one could use a mouse to select a command by using a pop-up menu, a fixed menu (a palette or command bar), multiple clicking, circling the desired command, or even writing the name of the command with the mouse. Software might even detect patterns of mouse use in the background, such as repeated “surfing” through menus, to automatically suggest commands or help topics (Horvitz, Breese, Heckerman, Hovel & Rommelse, 1998). The latter suggests a shift from the classical view of interaction as direct manipulation where the user is responsible for all actions and decisions, to one which uses **background sensing** techniques to allow technology to support the user with semi-automatic or implicit actions and services (Buxton, 1995a).

The Composition of Interaction Tasks

Early efforts in human-computer interaction sought to identify elemental tasks that appear repeatedly in human-computer dialogs. Foley, Wallace, and Chan proposed that user interface transactions are composed of the following elemental tasks:

Selection: Choosing objects from a set of alternatives

Position: Specifying a position within a range. This includes picking a screen coordinate with a pointing device.

Orient: Specifying an angle or three-dimensional orientation.

Path: Specifying a series of positions and/or orientations over time.

Quantify: Specifying an exact numeric value.

Text: Entry of symbolic data.

While these are commonly occurring tasks in many direct-manipulation interfaces, a problem with this approach is that the level of analysis at which one specifies “elemental” tasks is not well defined. For example, for *Position* tasks, a screen coordinate could be selected using a pointing device such as a mouse, but might be entered as a pair of numeric values (*Quantify*) using a pair of knobs (like an Etch-a-Sketch) where precision is paramount. But if these represent elemental tasks, why do we find that we must subdivide *Position* into a pair of *Quantify* subtasks for some devices but not for others?

Treating all tasks as hierarchies of sub-tasks, known as **compound tasks**, is one way to address this. With appropriate design, and by using technologies and interaction metaphors that parallel the way the user thinks about a task as closely as possible, the designer can phrase together a series of elemental tasks into a single **cognitive chunk**. For example, if the user’s task is to draw a rectangle, a device such as an Etch-a-Sketch is easier to use. For drawing a circle, a pen is far easier to use. Hence the choice of device influences the level at which the user is required to think about the individual actions that must be performed to achieve a goal. See (Buxton, 1986) for further discussion of this important concept.

The six elemental tasks enumerated above may be a complete list of “fundamental” low-level tasks that underlie most interaction with computers, but it could be argued that this list is not complete; for example, which of these six tasks does a fingerprint scanner support? Perhaps, if used for password replacement, it could be viewed as supporting the **Text** task; alternatively, one might add “Establishment of Identity” to the list. This points to a problem with the fundamental task approach. While identifying “elemental tasks” can be useful for thinking about interaction techniques in general, a problem with viewing tasks as assemblies of elemental tasks is that it typically only considers explicit input in the classical direct manipulation paradigm. Where do devices like cameras, microphones, and the fingerprint scanner discussed above fit in? These support higher-level data types and concepts (e.g. images, audio, and identity). Advances in technology will continue to yield new “elemental” inputs. However, these new technologies also may make increasing demands on systems to move from individual samples to synthesis of meaningful structure from the resulting data (Fitzmaurice, Balakrishnan & Kurtenbach, 1999).

Properties of Input Devices

The breadth of input devices and displays on the market today can be completely bewildering. Fortunately, there are a number of organizing properties and principles which can help to make sense of the design space and performance issues. First, we consider continuous, manually operated pointing devices (as opposed to discrete input mechanisms such as buttons or keyboards, or other devices not operated with the hand, which we will discuss briefly later). For

further insight readers may also wish to consult complete taxonomies of devices ((Buxton, 1983; Card, Mackinlay & Robertson, 1991)). As we shall see, however, it is nearly impossible to describe properties of input devices without reference to output—especially the resulting feedback on the screen—since after all input devices are only useful insofar as they support interaction techniques that allow the user to accomplish something.

Physical property sensed. Traditional pointing devices typically sense position, motion, or force. A tablet senses position, a mouse measures motion (i.e. change in position), and an isometric joystick senses force. An isometric joystick is a self-centering force sensing joystick such as the IBM TrackPoint (“eraser-head”) found on many laptops. For a rotary device, the corresponding properties are angle, change in angle, and torque. Position sensing devices are also known as **absolute input devices**, whereas motion sensing devices are **relative input devices**. An absolute device can fully support relative motion, since it can calculate changes to position, but a relative device cannot fully support absolute positioning, and in fact can only emulate “position” at all by introducing a cursor on the screen. Note that it is difficult to move the mouse cursor to a particular area of the screen (other than the edges) without looking at the screen, but with a tablet one can easily point to a region with the stylus using the kinesthetic sense (Balakrishnan & Hinckley, 1999), informally known as “muscle memory.”

Transfer function. A device, in combination with the host operating system, typically modifies its signals using a mathematical transformation that scales the data to provide smooth, efficient,

and intuitive operation. An appropriate mapping is a transfer function that matches the physical properties sensed by the input device. Appropriate mappings include force-to-velocity, position-to-position, and velocity-to-velocity functions. For example, an isometric joystick senses force; a nonlinear rate mapping transforms this into a velocity of cursor movement (Rutledge & Selker, 1990; Zhai & Milgram, 1993; Zhai, Smith & Selker, 1997). Ideally, the device should also be self-centering when using a rate mapping, with a spring return to the zero input value, so that the user can stop quickly by releasing the device. A common inappropriate mapping is calculating a speed of scrolling based on the position of the mouse cursor, such as extending a selected region by dragging the mouse close to the edge of the screen. The user has no feedback of when or to what extent scrolling will accelerate, and the resulting interaction can be hard to learn how to use and difficult to control.

A simple multiplicative transfer function is known as the device gain, which can also be described as a **control-to-display (C:D) ratio**, the ratio between the movement of the input device and the corresponding movement of the object it controls. For example, if a mouse (the device) must be moved 1 cm on the desk in order to move a cursor 2 cm on the screen (the display), the device has a 1:2 control-display ratio. However, on commercial pointing devices and operating systems, the gain is rarely constant¹; an *acceleration function* is often used to

¹ Direct input devices are an exception, since the C:D ratio is typically fixed at 1:1 (but see also Sears, A. & B. Shneiderman (1991). "High Precision Touchscreens: Design Strategies and Comparisons with a Mouse." International Journal of Man-Machine Studies 34(4): 593-613.).

modulate the gain depending on velocity. An acceleration function is a transfer function that exhibits an exponential relationship between velocity and gain. Experts believe the primary benefit of acceleration is to reduce the **footprint**, or the physical movement space, required by an input device (Jellinek & Card, 1990; Hinckley, Cutrell, Bathiche & Muss, 2001). One must also be very careful when studying the possible influence of gain settings on user performance: experts have criticized gain as a fundamental concept, since it confounds two separate concepts (device size and display size) in one arbitrary metric (Accot & Zhai, 2001). Furthermore, user performance may exhibit speed-accuracy tradeoffs, calling into question the assumption that there exists an “optimal” C:D ratio (MacKenzie, 1995).

Number of dimensions. Devices can measure one or more linear and angular dimensions. For example, a mouse measures two linear dimensions, a knob measures one angular dimension, and a six degree of freedom (6DOF) magnetic tracker measures three linear dimensions and three angular (for examples of 6DOF input and design issues, see (Ware & Jessome, 1988; Hinckley, Pausch, Goble & Kassell, 1994b; Green & Liang, 1994; Serra, Hern, Beng Choon & Poston, 1997)). If the number of dimensions required by the user’s interaction task does not match the number of dimensions provided by the input device, then special handling (e.g. interaction techniques that may require extra buttons, graphical widgets, mode switching, etc) will need to be introduced. This is particularly a concern for three-dimensional user interfaces and interaction (Zhai, 1998; Hinckley et al., 1994b). Numerous interaction techniques have been proposed to allow standard 2D pointing devices to control 3D positioning or orientation tasks (e.g. (Conner et

al., 1992), (Chen, Mountford & Sellen, 1988), (Bukowski & Sequin, 1995)). Well-designed interaction techniques using specialized multiple degree-of-freedom input devices can sometimes offer superior performance (Ware & Rose, 1999; Hinckley, Tullio, Pausch, Proffitt & Kassell, 1997), but may be ineffective for standard desktop tasks, so overall performance must be considered (Balakrishnan, Baudel, Kurtenbach & Fitzmaurice, 1997; Hinckley, Sinclair, Hanson, Szeliski & Conway, 1999).

Pointing speed and accuracy. The standard way to characterize pointing device performance employs the **Fitts' Law** paradigm (Fitts, 1954). Fitts' Law relates the *movement time* to point at a target, the *amplitude* of the movement (the distance to the target), and the *width* of the target (i.e., the precision requirement of the pointing movement). The movement time is proportional to the logarithm of the distance divided by the target width, with constant terms that vary from one device to another. While not emphasized in this chapter, Fitts' Law is the single most important quantitative analysis, testing, and prediction tool available to input research and device evaluation. For an excellent overview of its application to the problems of HCI, including use of Fitts' law to characterize bandwidth (a composite measure of both speed and accuracy), see (MacKenzie, 1992). For discussion of other accuracy metrics, see (MacKenzie, Kauppinen & Silfverberg, 2001). Recently the Fitts' Law testing paradigm has been proposed as an international standard for evaluating pointing devices (Douglas, Kirkpatrick & MacKenzie, 1999).

Recent years have seen a number of new insights and new applications for Fitts' Law. Fitts' Law was originally conceived in the context of rapid, aimed movements, but Fitts' law can also be applied to tasks such as scrolling (Hinckley et al., 2001), multi-scale navigation (Guiard, Buourgeois, Mottet & Beaudouin-Lafon, 2001), and crossing boundaries (Accot & Zhai, 2002). Researchers have also recently applied Fitts' Law to expanding targets that double in width as the user approaches them. Even if the expansion begins after the user has already covered 90% of the distance from a starting point, the expanding target can be selected as easily as if it had been fully expanded since the movement began (McGuffin & Balakrishnan, 2002); see also (Zhai, Conversy, Beaudouin-Lafon & Guiard, 2003). However, it remains unclear if this can be successfully applied to improve pointing performance for multiple targets that are closely packed together (as typically found in menus and tool palettes). For tasks that exhibit continuous speed-accuracy requirements, such as moving through a hierarchical menu, Fitts' Law cannot be applied, but researchers have recently formulated the Steering Law, which does address such tasks (Accot & Zhai, 1997; Accot & Zhai, 1999; Accot & Zhai, 2001).

Input Device States. To select a single point or region with an input device, users need a way to signal when they are selecting something versus when they are just moving over something to reach a desired target. The need for this fundamental signal of intention is often forgotten by researchers eager to explore new interaction modalities such as empty-handed pointing (e.g. using camera tracking or non-contact proximity sensing of hand position). The **three-state model** of input (Buxton, 1990b) generalizes the states sensed by input devices as *tracking*,

which causes the cursor to move, *dragging*, which allows selection of objects by clicking (as well as moving objects by clicking and dragging them), and *out of range*, which occurs when the device moves out of its physical tracking range (e.g. a mouse is lifted from the desk, or a stylus is removed from a tablet). Most pointing devices sense only two of these three states: for example, a mouse senses tracking and dragging, but a touchpad senses tracking and the out of range state. Hence, to fully simulate the functionality offered by mice, touchpads need special procedures, such as tapping to click, which are prone to inadvertent activation (e.g. touching the pad by accident causes a click (MacKenzie & Oniszczak, 1998)). For further discussion and examples, see (Buxton, 1990b) and (Hinckley, Czerwinski & Sinclair, 1998a).

Direct vs. indirect control. A mouse is an **indirect input device** (one must move the mouse to point to a spot on the screen); a touchscreen is a **direct input device** (the display surface is also the input surface). Direct devices raise several unique issues. Designers must consider the possibility of parallax error resulting from a gap between the input and display surfaces, reduced transmissivity of the screen introduced by a sensing layer, or occlusion of the display by the user's hands. Another issue is that touchscreens can support a cursor tracking state, or a dragging state, but not both; typically touchscreens move directly from the out of range state to the dragging state when the user touches the screen, with no intermediate cursor feedback (Buxton, 1990b). Techniques for touchscreen cursor feedback have been proposed, but typically require that selection occurs on lift-off (Sears & Shneiderman, 1991; Sears, Plaisant & Shneiderman, 1992). See also Pen Input in the Discussion of Common Pointing Devices later in this chapter.

Device acquisition time is The average time to pick up or put down an input device is known as **acquisition time** (or sometimes homing time). It is often assumed to be a significant factor for user performance, but the Fitts' Law bandwidth of a device tends to dominate acquisition time unless switching occurs frequently (Douglas & Mithal, 1994). However, one exception is stylus or pen-based input devices; pens are generally comparable to mice in general pointing performance (Accot & Zhai, 1999), or even superior for some high-precision tasks (Guiard et al., 2001), but these benefits can easily be negated by the much greater time it takes to switch between using a pen and using a keyboard.

Modeling of repetitive tasks. The keystroke-level model (KLM) is commonly used to model expert user performance in repetitive tasks such as text editing. The KLM includes standard operators that represent average times required for pointing with an input device, pressing a button, pauses for decision making, and device acquisition time, but the model does not account for errors or non-expert behaviors such as problem solving (Card, Moran & Newell, 1980). Good examples of research that apply the KLM include (Wang, Zhai & Su, 2001) and (MacKenzie & Soukoreff, 2002). Another approach called GOMS modeling is an extension of the KLM that can handle more complex cases (Olson & Olson, 1990), but many practitioners still use the KLM to evaluate input devices and low-level interaction techniques because of KLM's greater simplicity.

Hardware criteria. Various other characteristics can distinguish input devices, but are perhaps

less important in distinguishing the fundamental types of interaction techniques that can be supported. Engineering parameters of a device's performance such as sampling rate, resolution, accuracy, and linearity can all influence performance. **Latency** is the end-to-end delay between the user's physical movement, sensing this, and providing the ultimate system feedback to the user. Latency can be a devious problem as it is impossible to completely eliminate from system performance; latency of more than 75-100 milliseconds significantly impairs user performance for many interactive tasks (Robertson, Card & Mackinlay, 1989; MacKenzie & Ware, 1993). For vibrotactile or haptic feedback, users may be sensitive to much smaller latencies of just a few milliseconds (Cholewiak & Collins, 1991).

Discussion of Common Pointing Devices

Here, we briefly describe commonly available pointing devices and some issues that can arise with them in light of the properties discussed above.

Mouse: A mouse senses movement relative to a flat surface. Mice exhibit several properties that are well suited to the demands of desktop graphical interfaces (Balakrishnan et al., 1997).

The mouse is stable and does not fall over when released (unlike a stylus on a tablet). A mouse can also provide integrated buttons for selection, and since the force required to activate a mouse's buttons is orthogonal to the plane of movement, it helps minimize accidental clicking or interference with motion. Another subtle benefit is the possibility for users to employ a

combination of finger, hand, wrist, arm, and even shoulder muscles to span the range of tasks from short precise selections to large, ballistic movements (Zhai, Milgram & Buxton, 1996; Balakrishnan & MacKenzie, 1997). Finally, Fitts' law studies show that users can point with the mouse about as well as with the hand itself (Card, English & Burr, 1978).

Trackball: A trackball is like a mouse that has been turned upside-down, with a mechanical ball that rolls in place. The main advantage of trackballs is that they can be used on an inclined surface and they often require a smaller footprint than mice. They also employ different muscle groups, which some users find more comfortable. However, trackballs cannot accommodate buttons as easily as mice, so tasks which require moving the trackball while holding down a button can be awkward (MacKenzie, Sellen & Buxton, 1991).

Tablets: Most tablets sense the absolute position of a mechanical intermediary such as a stylus or puck on the tablet surface. A puck is a mouse that is used on a tablet; the only difference is that it senses absolute position and it cannot be used on a surface other than the tablet. Absolute mode is generally preferable for tasks such as tracing, digitizing, drawing, free-hand inking, and signature capture. Tablets which sense contact of the bare finger are known as touch tablets (Buxton, Hill & Rowley, 1985); touchpads are miniature touch tablets, as commonly found on portable computers (MacKenzie & Oniszczak, 1998). A touchscreen is a transparent touch-sensitive tablet mounted on a display, but demands different handling than a tablet (see direct vs. indirect above).

Pen Input. Pen-based input for mobile devices is an area of increasing practical concern. Pens effectively support activities such as inking, marking, and gestural input (see the section on Pen-Based Gestures and Hand Gesture Input below), but raise a number of problems when supporting graphical interfaces originally designed for mouse input. Pen input raises the concerns of direct input devices as described above. There is no way to see exactly what position will be selected before selecting it: pen contact with the screen directly enters the Dragging state of the three-state model (Buxton, 1990b). There is no true equivalent of a “hover” state for tool tips¹, nor an extra button for context menus. Pen dwell time on a target can be used to provide one of these two functions. For detecting double-tap, allow a longer interval between the taps (as compared to double-click on a mouse), and also allow a significant change to the screen position between taps. Finally, users often want to touch the screen of small devices using a bare finger, so applications should be designed to accommodate imprecise selections. Note that some pen-input devices, such as the Tablet PC, use an inductive sensing technology that can only sense contact from a specially instrumented stylus, and thus cannot be used as a touchscreen. However, this deficiency is made up for by the ability to track the pen when it is close to (but not touching) the screen, allowing support for a tracking state with cursor feedback (and hence tool tips as well).

¹ Tool tips are small explanatory labels or balloons that appear next to a button, icon, or other interface widget when the user holds the cursor still over that object. The “hover” state is detected when there is little or no cursor movement for a fixed timeout (many systems use a timeout of approximately one second).

Joysticks. There are many varieties of joysticks. As mentioned above, an isometric joystick senses force and returns to center when released. Because isometric joysticks can have a tiny footprint, they are often used when space is at a premium, allowing integration with a keyboard and hence rapid switching between typing and pointing (Rutledge & Selker, 1990; Douglas & Mithal, 1994). Isotonic joysticks sense the angle of deflection of the stick; they tend to move more than isometric joysticks, offering better feedback to the user. Such joysticks may or may not have a mechanical spring return to center. Some joysticks even include both force and position sensing, and other special features. For a helpful organization of the complex design space of joysticks, see (Lipscomb & Pique, 1993).

Alternative Means of Pointing. Researchers have explored using the feet (Pearson & Weiser, 1988), head tracking, and eye tracking as alternative approaches to pointing. Head tracking has much lower pointing bandwidth than the hands, and may require the neck to be held in an awkward fixed position, but has useful applications for intuitive coupling of head movements to virtual environments (Sutherland, 1968; Brooks, 1988) and interactive 3D graphics (Hix, Templeman & Jacob, 1995; Ware, Arthur & Booth, 1993). Eye movement-based input, properly used, can provide an unusually fast and natural means of communication, because we move our eyes rapidly and almost unconsciously. The human eye fixates visual targets within the fovea, which fundamentally limits the accuracy of eye gaze tracking to 1 degree of the field of view (Zhai, Morimoto & Ihde, 1999). Eye movements are subconscious and must be interpreted carefully to avoid annoying the user with unwanted responses to his actions, known as the Midas

touch problem (Jacob, 1991). Current eye tracking technology is expensive and has numerous technical limitations, confining its use thus far to research labs and disabled persons with few other options.

Feedback and Perception-Action Coupling

The ecological approach to human perception (Gibson, 1986) asserts that the organism, the environment, and the tasks the organism performs are inseparable and should not be studied in isolation. Hence perception and action are intimately linked in a single motor-visual feedback loop, and any separation of the two is an artificial one. The lesson for interaction design is that techniques must consider both the motor control (input) and feedback (output) aspects of the design and how they interact with one another.

From the technology perspective, one can consider feedback as passive or active. Active feedback is under computer control. This can be as simple as presenting a window on a display, or as sophisticated as simulating haptic contact forces with virtual objects when the user moves an input device. We will return to discuss active feedback techniques later in this chapter, when we discuss display technologies and techniques.

Passive feedback may come from sensations within the user's own body, as influenced by physical properties of the device, such as the shape, color, and feel of buttons when they are depressed. The industrial design of a device suggests the purpose and use of a device even before

a user touches it (Norman, 1990). Mechanical sounds and vibrations that result from using the device provide confirming feedback of the user's action. The shape of the device and the presence of landmarks can help users orient a device without having to look at it (Hinckley, Pausch, Proffitt & Kassell, 1998b). Proprioceptive and kinesthetic feedback are somewhat imprecise terms, often used interchangeably, that refer to sensations of body posture, motion, and muscle tension (Mackenzie & Iberall, 1994). These senses allow users to feel how they are moving an input device without looking at the device, and indeed without looking at the screen in some situations (Mine, Brooks & Sequin, 1997; Balakrishnan & Hinckley, 1999). This may be important when the user's attention is divided between multiple tasks and devices (Fitzmaurice & Buxton, 1997). Sellen et al. (Sellen, Kurtenbach & Buxton, 1992) report that muscular tension from depressing a foot pedal makes modes more salient to the user than purely visual feedback. Although all of these sensations are passive and not under the direct control of the computer, these examples nonetheless demonstrate that they are relevant to the design of devices, and interaction techniques can consider these qualities and attempt to leverage them.

User performance may be influenced by correspondences between input and output. Some correspondences are obvious, such as the need to present confirming visual feedback in response to the user's actions. Ideally, feedback should indicate the results of an operation before the user commits to it (e.g., highlighting a button or menu item when the cursor moves over it).

Kinesthetic correspondence and perceptual structure, described below, are less obvious.

Kinesthetic correspondence refers to the principle that graphical feedback on the screen should correspond to the direction that the user moves the input device, particularly when 3D rotation is involved (Britton, Lipscomb & Pique, 1978). Users can easily adapt to certain non-correspondences: when the user moves a mouse forward and back, the cursor actually moves up and down on the screen; if the user drags a scrollbar downward, the text on the screen scrolls upwards. With long periods of practice, users can adapt to almost anything (for example, for over 100 years psychologists have known of the phenomena of prism adaptation, where people can eventually adapt to wearing prisms that cause everything to look upside-down (Stratton, 1897)). However, one should not force users to adapt to a poor design.

Perceptual Structure. Researchers have also found that the interaction of the input dimensions of a device with the control dimensions of a task can exhibit perceptual structure. Jacob, Sibert, McFarlane & Mullen (Jacob, Sibert, McFarlane & Mullen, 1994) explored two input devices, a 3D position tracker with integral (x, y, z) input dimensions, and a standard 2D mouse, with (x, y) input separated from (z) input by holding down a mouse button. For selecting the position and size of a rectangle, the position tracker is most effective: here, the integral 3D position input matches the integral presentation of the feedback on the screen. But for selecting the position and grayscale color of a rectangle, the mouse is most effective: here, the user perceives the position and grayscale color of the rectangle as separate quantities, and can more easily perform the task when the input device provides separate controls. Hence neither a 3D integral input nor a 2D (x, y) plus 1D (z) input is uniformly superior; the better performance results when the task and

device are both integral or both separable.

Keyboards, Text Entry, and Command Input

For over a century, keyboards and typewriters have endured as the mechanism of choice for text entry. The resiliency of the keyboard, in an era of unprecedented technological change, is the result of how keyboards complement human skills, and may make keyboards difficult to supplant with new input devices or technologies. We summarize some general issues surrounding text entry below, with a focus on mechanical keyboards; see also (Lewis, Potosnak & Magyar, 1997).

Skill Acquisition and Skill Transfer: Procedural memory is a specific type of memory that encodes repetitive motor acts. Once an activity is encoded in procedural memory, it requires little conscious effort to perform (Anderson, 1980). Because procedural memory automates the physical act of text entry, touch-typists can rapidly type words without interfering with the mental composition of text. The process of encoding an activity in procedural memory can be formalized as the **power law of practice**: $T = aP^b$, where T is the time to perform the task, P is the amount of practice, and a and b are constants that fit the curve to observed data. This suggests that changing the keyboard can have a high re-learning cost. However, a change to the keyboard can succeed if it does not interfere with existing skills, or allows a significant transfer of skill. For example, some ergonomic keyboards preserve the basic key layout, but alter the typing pose to help maintain neutral postures (Honan, Serina, Tal & Rempel, 1995; Marklin & Simoneau, 1996), whereas the Dvorak key layout may have some small performance advantages,

but has not found wide adoption due to high retraining costs (Lewis et al., 1997).

Eyes-Free Operation: With practice, users can memorize the location of commonly used keys relative to the home position of the two hands, allowing typing with little or no visual attention (Lewis et al., 1997). By contrast, soft keyboards (small on-screen virtual keyboards found on many hand-held devices) require nearly constant visual monitoring, resulting in diversion of attention from one's work. Furthermore, with stylus-driven soft keyboards, the user can only strike one key at a time. Thus the design issues for soft keyboards differ tremendously from mechanical keyboards (Zhai, Hunter & Smith, 2000).

Tactile Feedback: On a mechanical keyboard users can feel the edges and gaps between the keys, and the keys have an activation force profile that provides feedback of the key strike. In the absence of such feedback, as on touchscreen keyboards (Sears, 1993), performance may suffer and users may not be able to achieve eyes-free performance (Lewis et al., 1997).

Combined Text, Command, and Navigation Input. Finally, it is easy to forget that keyboards provide many secondary command and control actions in addition to pure text entry, such as power keys and navigation keys (Enter, Home/End, Delete, Backspace, Tab, Esc, Page Up/Down, Arrow Keys, etc.), chord key combinations (such as Ctrl+C for Copy) for frequently used commands, and function keys for miscellaneous functions defined by the current application. Without these keys, frequent interleaving of mouse and keyboard activity may be

required to perform these secondary functions.

Ergonomic Issues. Many modern information workers suffer from repetitive strain injury (RSI). Researchers have identified many risk factors for such injuries, such as working under stress or taking inadequate rest breaks. People often casually associate these problems with keyboards, but the potential for RSI is common to many manually operated tools and repetitive activities (Putz-Anderson, 1988). Researchers have advocated themes for ergonomic design of keyboards and other devices (Pekelney & Chu, 1995), including reducing repetition, minimizing force required to hold and move the device or press its buttons, avoiding sharp edges that put pressure on the soft tissues of the hand, and designing for natural and neutral postures of the user's hands and wrists (Honan et al., 1995; Marklin, Simoneau & Monroe, 1997). Communicating a clear orientation for gripping and moving the device through its industrial design also may help to discourage inappropriate, ergonomically unsound grips.

Other text entry mechanisms. One-handed keyboards can be implemented using simultaneous depression of multiple keys; such *chord keyboards* can sometimes allow one to achieve high peak performance (e.g. court stenographers), but take much longer to learn how to use (Noyes, 1983; Mathias, MacKenzie & Buxton, 1996; Buxton, 1990a). They are often used in conjunction with wearable computers (Smailagic & Siewiorek, 1996) to keep the hands free as much as possible (but see also Voice and Speech below). With complex written languages, such as Chinese and Japanese, key chording and multiple stages of selection and disambiguation are

currently necessary for keyboard-based text entry (Wang et al., 2001). Handwriting and character recognition may ultimately provide a more natural solution, but for Roman languages handwriting (even on paper, with no recognition involved) is much slower than skilled keyboard use. To provide reliable stylus-driven text input, some systems have adopted unistroke (single-stroke) gestural “alphabets” (Goldberg & Richardson, 1993) that reduce the demands on recognition technology, while remaining relatively easy for users to learn (MacKenzie & Zhang, 1997). However, small “two thumb” keyboards (MacKenzie & Soukoreff, 2002) or fold-away peripheral keyboards are becoming increasingly popular for mobile devices. Dictation using continuous speech recognition is available on the market today, but the technology still has a long way to go; a recent study found that the corrected words-per-minute rate of text entry using a mouse and keyboard are about twice as fast as dictation input (Karat, Halverson, Horn & Karat, 1999). We further discuss speech interaction in the next section.

Modalities of Interaction

Here, we briefly review a number of general strategies and input modalities which have been explored by researchers. These approaches generally transcend a specific type of input device, but rather span a range of devices and applications.

Voice and Speech

Carrying on a full conversation with a computer as one might do with another person is well

beyond the state of the art today and, even if possible, may be a naive goal. Yet even without understanding the content of the speech, computers can digitize, store, edit, and replay segments of speech to augment human-human communication (Arons, 1993; Stifelman, 1996; Buxton, 1995b). Conventional voice mail and the availability of MP3 music files on the web are simple examples of this. Computers can also infer information about the user's activity from ambient audio, such as determining if the user is present, or perhaps engaging in a conversation with a colleague, allowing more timely delivery of information, or suppression of notifications that may interrupt the user (Schmandt, Marmasse, Marti, Sawhney & Wheeler, 2000; Sawhney & Schmandt, 2000; Horvitz, Jacobs & Hovel, 1999).

Understanding speech as input has been a long-standing area of research. While progress is being made, it is slower than optimists originally predicted, and daunting unsolved problems remain. For limited vocabulary applications with native English speakers, speech recognition can excel at recognizing words that occur in the vocabulary. Error rates can increase substantially when users employ words that are out-of-vocabulary (i.e. words the computer is not "listening" for), when the complexity of the grammar of possible phrases increases, or when the microphone is not a high-quality close-talk headset. Even if the computer could recognize all of the user's words, the problem of understanding natural language is a significant and unsolved one. It can be avoided by using an artificial language of special commands or even a fairly restricted subset of natural language. But, given the current state of the art, the closer the user moves toward full unrestricted natural language, the more difficulties will be encountered.

Pen-Based Gestures, Hand Gesture Input, and Multimodal Input

Pen-based gestures can indicate commands, such as crossing out a word to delete it, or circling a paragraph and drawing an arrow to move it. Such gestures support cognitive chunking by integrating command selection with specification of the command's scope (Buxton, Fiume, Hill, Lee & Woo, 1983; Kurtenbach & Buxton, 1991). Marking menus use directional pen motion to provide extremely rapid menu selection (Kurtenbach & Buxton, 1993; Kurtenbach, Sellen & Buxton, 1993). With pen-based interfaces, designers often face a difficult design trade-off between treating the user's marks as ink that is not subject to interpretation, versus providing pen-based input that treats the ink as a potential command (Kramer, 1994; Moran, Chiu & van Melle, 1997; Mynatt, Igarashi, Edwards & LaMarca, 1999). Pen input, via sketching, can be used to define 3D objects (Zelevnik, Herndon & Hughes, 1996; Igarashi, Matsuoka & Tanaka, 1999). Researchers have also explored multimodal pen and voice input; this is a powerful combination as pen and voice have complementary strengths and weaknesses, and can disambiguate one another (Cohen et al., 1997; Cohen & Sullivan, 1989; Oviatt, 1997).

Input using hand gestures represents another, less well understood area of inquiry. There are many human behaviors involving hand movements and gestures, but few have been thoroughly explored for human-computer interaction. Cadoz (Cadoz, 1994) broadly categorizes hand gestures as semiotic, ergotic, or epistemic. Semiotic gestures are those used to communicate meaningful information, such as "thumbs up." Ergotic gestures are those used to manipulate

physical objects. Epistemic gestures are exploratory movements to acquire haptic or tactile information; see also (Kirsh, 1995; Kirsh & Maglio, 1994). Most research in hand gesture recognition focuses on empty-handed semiotic gestures (Cassell, 2003), which Rime and Schiaratura (Rime & Schiaratura, 1991) further classify as symbolic (conventional symbolic gestures such as “OK”), deictic (pointing to fill in a semantic frame, analogous to deixis in natural language), iconic (illustrating a spatial relationship), or pantomimic (mimicking an invisible tool, such as pretending to swing a golf club).

Command input using recognition-based techniques raises a number of unique challenges (Bellotti et al., 2002). In particular, with most forms of gestural input, errors of user intent and errors of computer interpretation seem inevitable. Deictic gesture in particular has received much attention, with several efforts using pointing (typically captured using instrumented gloves or camera-based recognition) to interact with “intelligent” environments (Baudel & Beaudouin-Lafon, 1993; Maes, Darrell, Blumberg & Pentland, 1996; Freeman & Weissman, 1995; Jojic, Brumitt, Meyers & Harris, 2000) as well as deictic gesture in combination with speech recognition (Bolt, 1980; Hauptmann, 1989; Lucente, Zwart & George, 1998; Wilson & Shafer, 2003). However, there is more to the field than empty-handed semiotic gesture. Recent exploration of tangible interaction techniques (Ishii & Ullmer, 1997) and efforts to sense movements and handling of sensor-enhanced mobile devices represent examples of sensing manipulation (that is, ergotic gestures) (Hinckley, Pierce, Horvitz & Sinclair, 2003; Hinckley, Pierce, Sinclair & Horvitz, 2000; Harrison, Fishkin, Gujar, Mochon & Want, 1998).

Bimanual Input

Aside from touch typing, most of the devices and modes of operation discussed thus far and in use today involve only one hand at a time. But people use both hands in a wide variety of the activities associated with daily life. For example, when writing, a right-hander writes with the pen in the right hand, but the left hand also plays a crucial and distinct role. It holds the paper and orients it to a comfortable angle that suits the right hand. In fact, during many skilled manipulative tasks, Guiard observed that the hands take on asymmetric, complementary roles (Guiard, 1987): for right-handers, the role of the left hand precedes the right (the left hand first positions the paper), the left hand sets the frame of reference for the action of the right hand (the left hand orients the paper), and the left hand performs infrequent, large-scale movements compared to the frequent, small-scale movements of the right hand (writing with the pen). Most applications for bimanual input to computers are characterized by asymmetric roles of the hands, including compound navigation/selection tasks such as scrolling a web page and then clicking on a link (Buxton & Myers, 1986), command selection using the nonpreferred hand (Bier, Stone, Pier, Buxton & DeRose, 1993; Kabbash, Buxton & Sellen, 1994), as well as navigation, virtual camera control, and object manipulation in three-dimensional user interfaces (Kurtenbach, Fitzmaurice, Baudel & Buxton, 1997; Balakrishnan & Kurtenbach, 1999; Hinckley et al., 1998b). Researchers have also applied this approach to keyboard design (MacKenzie & Guiard, 2001; McLoone, Hinckley & Cutrell, 2003). For some tasks, such as banging together a pair of cymbals, the hands may take on symmetric roles; for further discussion of bimanual symmetric

tasks, see (Guiard, 1987) and (Balakrishnan & Hinckley, 2000).

Passive Measurement: Interaction in the Background

Not all interaction with computers need consist of explicit, intentionally communicated commands. Think about walking into a grocery store with automatic doors. You approach the building; the doors sense this motion and open for you. No explicit communication has occurred, yet a computer has used your action of walking towards the store as an “input” to decide when to open the door. Intentional, explicit interaction takes place in the *foreground*, while implicitly sensed interaction takes place in the *background*, behind the fore of the user’s attention (Buxton, 1995a). **Background sensing techniques** will be a major emphasis of future research in automation and sensing systems as users become increasingly mobile and become saturated with information from many sources. Researchers are currently exploring ways of providing context awareness through location sensing, ambient sensing of light, temperature, and other environmental qualities, movement and handling of devices, detecting the identity of the user and physical objects in the environment, and possibly even physiological measures such as heart-rate variability. This type of information potentially can allow technology to interpret the context of a situation and respond more appropriately (Schilit, Adams & Want, 1994; Schmidt, 1999; Dey, Abowd & Salber, 2001; Hinckley et al., 2003). However, like other recognition-based technologies, there is a risk of errors of user intent or computer interpretation: returning to the automatic door example, if you walk by (parallel to) the doors, they may sense your motion and open even if you have no intention of entering the building.

Background interaction can also be applied to explicit input streams through passive behavioral measurements, such as observation of typing speed, manner of moving the cursor, sequence and timing of commands activated in a graphical interface (Horvitz et al., 1998), and other patterns of use. For example, a carefully designed user interface could make intelligent use of such information to modify its dialogue with the user, based on inferences about the user's alertness or expertise. These measures do not require additional input devices, but rather gleaning of additional, typically neglected, information from the existing input stream. These are sometimes known as intelligent or adaptive user interfaces, but mundane examples also exist. For example, cursor control using the mouse or scrolling using a wheel can be optimized by modifying the device response depending on the velocity of movement (Jelinek & Card, 1990; Hinckley et al., 2001).

We must acknowledge the potential for misuse or abuse of information collected in the background. Users should always be made aware of what information is or may potentially be observed as part of a human-computer dialogue. Users should have control and the ability to block any information that they want to remain private (Nguyen & Mynatt, 2001).

Displays and Perception

We now turn our attention to focus on the fundamental properties of displays and techniques for

effective use of displays. We focus on visual displays and visual human perception, since these represent the vast majority of displays, but we also discuss feedback through the haptic and audio channels.

Properties of Displays and Human Visual Perception

Display requirements, such as resolution in time and space, derive from the properties of human vision. Thus we begin with the basic issues relating to display brightness, uniformity, and spatial and temporal resolution.

Dynamic Range. The human eye has an enormous dynamic range. The amount of light reflected from surfaces on a bright day at the beach is about five orders of magnitude higher than the amount available under dim lamp lights. Yet the shapes, layouts, and colors of objects look nearly identical to the human eye across much of this range. Most displays in common use are self-luminous cathode ray tubes (CRTs) or back lit liquid crystal displays (LCDs). The best of these devices have a dynamic range (the ratio between the maximum and minimum values produced) of a little more than two orders of magnitude. In practice, under typical room lighting conditions, 15-40% of the light reaching the user's eye is actually ambient room light reflected by the front surface of the phosphors, or off of the screen surface. This means that the effective dynamic range of most devices, unless viewed in dark rooms, is no better than three or four to one. Fortunately the human eye can tolerate extreme variation in the overall level of illumination, as well as the amount of contrast produced by the display.

Spatial Frequency. The ability of the human visual system to resolve fine targets is known as **visual acuity**. A standard way of measuring visual acuity is to determine how fine a sinusoidal striped pattern can be discriminated from a uniform gray. Humans are capable of perceiving targets as fine as 50-60 cycles/degree of visual angle when the pattern is of very high contrast. Fig. 1 illustrates the spatial sensitivity of the human eye as a function of spatial frequency. Specifically, it illustrates the degree of contrast required for sinusoidal gratings of different spatial frequencies to be perceived. The function has an inverted U shape with a peak at about 2 cycles/degree of visual angle. This means that 5-mm stripes at arms length are optimally visible. The falloff at low spatial frequencies indicates that the human visual system is insensitive to gradual changes in overall screen luminance. Indeed most CRTs have a brightness falloff toward the edges of as much as 20%, which we barely notice. This non-uniformity is even more pronounced with rear projection systems, due to the construction of screens that project light primarily in a forward direction. This is called the **screen gain**; a gain of 3.0 means that three times as much light is transmitted in the straight through direction compared to a perfect **Lambertian** diffuser. At other angles, less light is transmitted so that at a 45° off-axis viewing angle, only half as much light may be available compared to a perfect diffuser. Screen gain is also available with front projection with similar non-uniformities as a consequence, although the use of curved screens can compensate to some extent.

Fig. 1 Spatial contrast sensitivity function of the human visual system. There is a falloff in sensitivity to both detailed patterns (high-spatial frequencies) and to gradually changing gray values (low-spatial frequencies).

Spatial Resolution. The receptors in the human eye have a visual angle of about 0.8 seconds of arc. Modern displays provide approximately 40 pixels/cm. A simple calculation reveals that at about a 50-cm viewing distance, pixels will subtend about 1.5 seconds of arc, about two times the size of cone receptors in the center of vision. Viewed from 100 cm such a screen has pixels that will be imaged on the retina at about the same size as the receptors. This might suggest that we are in reach of the perfect display in terms of spatial resolution; such a screen would require approximately 80 pixels/cm at normal viewing distances. However, under some conditions the human visual system is capable of producing **superacuties** that imply resolution better than the receptor size. For example, during fusion in stereo vision, disparities smaller than 5 seconds of arc can be detected (Westheimer, 1979); see also (Ware, 2000) for a discussion of stereopsis and stereo displays aimed at the practitioner. Another example of superacuity is known as aliasing, resulting from the division of the screen into discrete pixels; for example, a line on the display that is almost (but not quite) horizontal may exhibit a jagged “stairstep” pattern that is very noticeable and unsatisfying. This effect can be diminished by **antialiasing**, which computes pixel color values that are averages of all the different objects that contribute to the pixel, weighted by the percentage of the pixel they cover. Similar techniques can be applied to improve the appearance of text, particularly on LCD screens, where individual red, green, and blue display

elements can be used for sub-pixel antialiasing (Betrissey et al., 2000; Platt, 2000).

Temporal Resolution, Refresh, and Update Rates. The **flicker fusion frequency** represents the least rapidly flickering light that the human eye does not perceive as steady. Flicker fusion frequency typically occurs around 50 Hz for a light that turns completely on and off (Wyszecki & Styles, 1982). In discussing the performance of monitors it is important to differentiate the **refresh rate** and the **update rate**. The refresh rate is the rate at which a screen is redrawn and it is typically constant (values of 60 Hz up to 120 Hz are common). By contrast, the update rate is the rate at which the system software updates the output to be refreshed. Ideally, this should occur at or above the refresh rate, but with increasingly demanding applications and complex data sets, this may not be possible. A rule of thumb states that a 10 Hz update rate is a minimum for smooth animation (Robertson et al., 1989). Motion blur (also known as temporal antialiasing) techniques can be applied to reduce the jerky effects resulting from low frame rates (Cook, 1986).

Color Vision and Color Displays

The single most important fact relating to color displays is that human color vision is trichromatic; our eyes contain three receptors sensitive to different wavelengths. For this reason it is possible to generate nearly all perceptible colors using only three sets of lights or printing inks. However, it is much more difficult to exactly specify colors using inks than using lights because, whereas lights can be treated as a simple vector space, inks interact in complex

nonlinear ways.

Luminance, Color Specification, and Color Gamut

Luminance is the standard term for specifying brightness, i.e. how much light is emitted by a self-luminous display. The luminance system in human vision gives us most of our information about the shape and layout of objects in space. The international standard for color measurement is the CIE (Commission International de L'Eclairage) standard. The central function in Fig. 3 is the CIE $V(\lambda)$ function, which represents the amount that light of different wavelengths contributes to the overall sensation of brightness. As this curve demonstrates, short wavelengths (blue) and long wavelengths (red) contribute much less than green wavelengths to the sensation of brightness. The CIE tristimulus functions, also shown in Fig. 3, are a set of color matching functions that represent the color vision of a typical person. Humans are most sensitive to the green wavelengths around 560 nm. Specifying luminance, or specifying a color in CIE tristimulus values, are complex technical topics; for further discussion, see (Ware, 2000) and (Wyszecki & Styles, 1982).

Fig. 2 The CIE tristimulus functions. These are used to represent the standard observer in colorimetry. Short wavelengths at the left-hand side appear blue, in the middle they are green, and to the right they are red. Humans are most sensitive to the green wavelengths around 560 nm.

A chromaticity diagram can be used to map out all possible colors perceptible to the human eye, as illustrated in Fig. 3. The pure spectral hues are given around the boundary of this diagram in nanometers (10^{-9} m). While the spacing of colors in tristimulus coordinates and on the chromaticity diagram is not perceptually uniform, **uniform color spaces** exist that produce a space in which equal metric distances are closer to matching equal perceptual differences (Wyszecki & Styles, 1982). For example, this can be useful to produce color sequences in map displays (Robertson, 1988).

The gamut of all possible colors is the dark gray region of the chromaticity diagram, with pure hues at the edge and neutral tones in the center. The triangular region represents the gamut achievable by a particular color monitor, determined by the colors of the phosphors given at the corners of the triangle. Every color within this triangular region is achievable, and every color outside of the triangle is not. This diagram nicely illustrates the tradeoff faced by the designer of color displays. A phosphor that produces a very narrow wavelength band will have chromaticity coordinates close to the pure spectral colors, and this will produce more saturated colors (thus enlarging the triangle). However, this narrow band also means that little light is produced.

The irregular shape inside the triangle illustrates the gamut of colors obtainable using printing inks. Notice that this set of colors is still smaller, causing difficulties when we try to obtain a hard copy reproduction of the colors on the monitor. Because the eye is relatively insensitive to overall color shifts and overall contrast changes, we can take the gamut from one device and map

it into the gamut of the printing inks (or some other device) by compressing and translating it. This is known as gamut mapping, a process designed to preserve the overall color relationships while effectively using the range of a device (Stone, Cowan & Beatty, 1988). However, it should be noted that the original colors will be lost in this process and that after a succession of gamut mappings colors may become distorted from their original values.

Fig. 3 A CIE chromaticity diagram with a monitor gamut and a printing ink gamut superimposed. The range of available colors with color printing is smaller than that available with a monitor, and both fall short of providing the full range of color that can be seen.

A process known as chromatic adaptation occurs in the human eye receptors and in the early stages of visual processing: for example, we hardly notice that daylight is much bluer than the yellow cast of tungsten light produced from ordinary light bulbs. The CIE standard does not account for chromatic adaptation, nor does it account for color contrast (colors appear differently to the human eye depending on the surrounding visual field). The practical implication is that we can get by with color monitors and printers that are grossly out of calibration. However, accurate color is essential in some applications. It is possible to precisely calibrate a color device so that the particular inputs required to produce a color may be specified in CIE tristimulus values. For monitor calibration, see (Cowan, 1983); for calibrating print devices, see Stone et al. (Stone et al., 1988). It is also possible to correct for the nonlinear response of CRT displays, a process

known as **gamma correction**, but keep in mind that CRT designers intentionally insert this nonlinearity to match the human eye's sensitivity to relative changes in light intensity. If one desires a set of perceptually equal gray steps, it is usually best to omit gamma correction. See (Ware, 2000) for further discussion.

Information Visualization

Researchers and practitioners have become increasingly interested in communicating large quantities of information quickly and clearly by leveraging the tremendous capabilities of the human visual system, a field known as information visualization. Thanks to advances in computer graphics hardware and algorithms, virtually all new desktop machines available today have sophisticated full-color displays with transparency and texture mapping for complex two-dimensional or three-dimensional scenes, and it now seems inevitable these capabilities will become commonplace on laptop computers, and ultimately even on handheld devices.

General Issues in Information Coding

The greatest challenge in developing guidelines for information coding is that there are usually effective alternatives, such as color, shape, size, texture, blinking, orientation, and gray value. Although a number of studies compare one or more coding methods separately, or in combination, there are so many interactions between the task and the complexity of the display that guidelines based on science are not generally practical. However, Tufte provides excellent

guidelines for information coding from an aesthetic perspective (Tufte, 1983; Tufte, 1990; Tufte, 1997). For further discussion, examples, and case studies, see also (Ware, 2000) and (Card, Mackinlay & Shneiderman, 1999).

A theoretical concept known as **preattentive processing** has interesting implications for whether or not the coding used can be processed in parallel by the visual system. The fact that certain coding schemes are processed faster than others is called the popout phenomenon, and this is thought to be due to early preattentive processing by the visual system. Thus, for example, the shape of the word “bold” is not processed preattentively, and it will be necessary to scan this entire page to determine how many times the word appears. However, if all of the instances of the word **bold** are emphasized, they pop out at the viewer. This is true as long as there are not too many other emphasized words on the same page: if there are less than seven or so instances they can be processed at a single glance. Preattentive processing is done for color, brightness, certain aspects of texture, stereo disparities, and object orientation and size. Codes that are preattentively discriminable are very useful if rapid search for information is desired (Triesman, 1985). The following visual attributes are known to be preattentive codes and, therefore, useful in differentiating information belonging to different classes:

- Color: use no more than 10 different colors for labeling purposes.
- Orientation: use no more than 10 orientations.
- Blink coding: use no more than 2 blink rates.
- Texture granularity: use no more than 5 grain sizes.

- Stereo depth: the number of depths that can be effectively coded is not known.
- Motion: objects moving out of phase with one another are perceptually grouped. The number of usable phases is not known.

However, coding multiple dimensions by combining different popout cues is not necessarily effective (Ware, 2000).

Color Information Coding

When considering information display, one of the most important distinctions is between chromatic and luminance information, because these are treated quite differently in human perception. Gray scales are not perceived in the same way as rainbow colored scales. A purely chromatic difference is one where two colors of identical luminance, such as red and green, are placed adjacent to one another. Research has shown that we are insensitive to a variety of information if it is presented through purely chromatic changes. This includes shape perception, stereo depth information, shape from shading, and motion. However, chromatic information helps us classify the material properties of objects. A number of practical implications arise from the differences in the way luminance and chromatic information is processed in human vision:

- Our spatial sensitivity is lower for chromatic information, allowing image compression techniques to transmit less information about hue relative to luminance.
- To make text visible it is important to make sure that there is a luminance difference between the color of the text and the color of the background . If the background may vary, it is a good idea to put a contrasting border around the letters (e.g. (Harrison &

Vicente, 1996)).

- When spatial layout is shown either through a stereo display or through motion cues, ensure adequate luminance contrast.
- When fine detail must be shown, for example, with fine lines in a diagram, ensure that there is adequate luminance contrast with the background.
- Chromatic codes are useful for labeling objects belonging to similar classes.
- Color (both chromatic and gray scale) can be used as a quantitative code, such as on maps, where it commonly encodes height and depth. However, simultaneous contrast effects can change the appearance of a patch of color depending on the surrounding colors; careful selection of colors can minimize this (Ware, 1988).

A number of empirical studies have shown color coding to be an effective way of identifying information. It is also effective if used in combination with other cues such as shape. For example, users may respond to targets faster if they can be identified by both shape and color differences (for useful reviews, see (Christ, 1975), (Stokes, Wickens & Kite, 1990) and (Silverstein, 1977)). Color codes are also useful in the perceptual grouping of objects. Thus, the relationship between a set of different screen objects can be made more apparent by giving them all the same color. However, it is also the case that only a limited number of color codes can be used effectively. The use of more than about 10 will cause the color categories to become blurred. In general, there are complex relationships between the type of symbols displayed, (e.g., point, line, area, or text), the luminance of the display, the luminance and color of the

background, and the luminance and color of the symbol (Spiker, Rogers & Cicinelli, 1985).

Integrated Control/Display Objects

When the purpose of a display is to allow a user to integrate diverse pieces of information, it may make sense to integrate the information into a single visual object or glyph (Wickens, 1992). For example, if the purpose is to represent a pump, the liquid temperature could be shown by changing the color of the pump, the capacity could be shown by the overall size of the pump, and the output pressure might be represented by the changing height of a bar attached to the output pipe, rather than a set of individual dials showing these attributes separately. However, perceptual distortions can result from an ill-chosen display mapping, and the object display may introduce visual clutter: if there are 50 pumps to control, then the outlines of all the pumps may interfere with the data of interest (Tufte, 1983). In object displays, input and output can be integrated in a manner analogous to widgets such as the scroll bar, or even more directly by having input devices that resemble the physical object being handled, known as a prop (Hinckley, Pausch, Goble & Kassell, 1994a). For some good examples of the linking of output and input see (Ahlberg & Shneiderman, 1994) as well as (Ishii & Ullmer, 1997).

This style of presentation and interaction can be especially relevant for telepresence or augmented reality applications, where the user needs to interact with actual physical objects which have attributes that must be viewed and controlled (Tani, Yamaashi, Tanikoshi, Futakawa & Tanifuji, 1992; Feiner, Macintyre & Seligmann, 1993). For more abstract data representation

tasks choosing the color, size, orientation, or texture to represent a particular data attribute may be difficult, and there seem to be practical limits on the number of attributes that one can encode simultaneously. Thus, object displays must usually be custom designed for each different display problem. In general, this means that the display and controls should somehow match the user's cognitive model of the task (Norman, 1990; Cole, 1986).

Three-Dimensional Graphics and Virtual Reality

Much research in three-dimensional information visualization and **virtual reality** is motivated by the observation that humans naturally operate in physical space, and can intuitively move about and remember where things are (an ability known as spatial memory). However, translating these potential benefits to artificially generated graphical environments is difficult because of limitations in display and interaction technologies. Virtual environments research pushed this to the limit by totally immersing the user in an artificial world of graphics, but this comes at the cost of visibility and awareness of colleagues and objects in the real world. This has led to research in so-called **fish tank virtual reality** displays by using a head tracking system in conjunction with a stereo display (Deering, 1992; Ware et al., 1993) or a mirrored setup, which allows superimposition of graphics onto the volume where the users hand's are located (Schmandt, 1983; Serra et al., 1997). However, much of our ability to navigate without becoming lost depends upon the vestibular system and spatial updating as we physically turn our bodies, neither of which are engaged with stationary displays (Loomis, Klatzky, Golledge & Philbeck, 1999; Chance, Gaunet, Beall & Loomis, 1998). For further discussion of navigation in

virtual environments, see (Darken & Sibert, 1995) and (Darken & Sibert, 1993); for application of spatial memory to three-dimensional environments, see (Robertson et al., 1998) and (Robertson et al., 1999).

Augmented Reality

Augmented reality superimposes information on the surrounding environment rather than blocking it out. For example, the user may wear a semitransparent display that has the effect of projecting labels and diagrams onto objects in the real world. It has been suggested that this may be useful for training people to use complex systems, or for fault diagnosis. For example, when repairing an aircraft engine the names and functions of parts could be made to appear superimposed on the parts seen through the display together with a maintenance record if desired (Caudell & Mizell, 1992; Feiner et al., 1993). The computer must obtain a detailed model of the environment, otherwise it is not possible to match the synthetic objects with the real ones. Even with this information, correct registration of computer graphics with the physical environment is an extremely difficult technical problem due to measurement error and system latency. This technology has been applied to heads-up displays for fighter aircraft, with semi-transparent information about flight paths and various threats in the environment projected on the screen in front of the pilot (Stokes et al., 1990), as well as digitally augmented desk surfaces (Wellner, 1993).

Scale in Displays

It is important to consider the full range of scale for display devices and form-factors which may embody an interaction task. Computer displays increasingly span orders of magnitude in size and available computational resources, from watches, hand-held PDA's (personal data assistants), tablet computers, desktop computers, all the way up to multiple-monitor and wall-size displays. A technique that works well on a desktop computer, such as a pull-down menu, may be awkward on a small handheld device, or even unusable on a wall-size display (where the top of the display may not even be within the user's reach). Each class of device seems to raise unique challenges, and the best approach may ultimately be to design special-purpose, appliance-like devices (see (Want & Borriello, 2000) for a survey) that suit specific purposes.

Small Displays

Users increasingly want to do more and more on handheld devices, mobile phones, pagers, and watches that offer less and less screen real estate. Researchers have investigated various strategies for conserving screen real estate. Transparent overlays allow divided attention between foreground and background layers (Harrison, Ishii, Vicente & Buxton, 1995; Harrison, Kurtenbach & Vicente, 1995; Harrison & Vicente, 1996; Kamba, Elson, Harpold, Stamper & Sukaviriya, 1996), but some degree of interference seems inevitable. This can be combined with sensing which elements of an interface are being used, such as presenting widgets on the screen only when the user is touching a pointing device (Hinckley & Sinclair, 1999). Researchers have also experimented with replacing graphical interfaces with *graspable* interfaces that respond to

tilting, movement, and physical gestures that do not need constant on-screen representations (Fitzmaurice, Ishii & Buxton, 1995; Rekimoto, 1996; Harrison et al., 1998; Hinckley et al., 2000). Much research in focus plus context techniques, including fisheye magnification (Bederson, 2000) and zooming metaphors (Perlin & Fox, 1993; Bederson et al., 1996; Smith & Taivalsaari, 1999), has also been motivated by providing more space than the boundaries of the physical screen can provide. Researchers have started to identify principles and quantitative models to analyze the tradeoffs between multiple views and zooming techniques (Baudisch, Good, Bellotti & Schraedley, 2002; Plumlee & Ware, 2002). There has been considerable effort devoted to supporting web browsing in extremely limited screen space (Buyukkokten, Garcia-Molina & Paepcke, 2001; Jones, Marsden, Mohd-Nasir, Boone & Buchanan, 1999; Trevor, Hilbert, Schilit & Koh, 2001).

Multiple Displays

Researchers have recently recognized that some very interesting design issues arise when multiple displays are considered, rather than the traditional single display of desktop computers. Having multiple monitors for a single computer is not like having one large display (Grudin, 2001). Users employ the boundary between displays to partition their tasks, with one monitor being reserved for a primary task, and other monitors being used for secondary tasks. Secondary tasks may support the primary task (e.g. reference material, help files, or floating tool palettes), may provide peripheral awareness of ongoing events (such as an e-mail client), or may provide other background information (to-do lists, calendars, etc.). Switching between applications has a

small time penalty (incurred once to switch, and again to return), and perhaps more importantly, it may distract the user or force the user to remember information while switching between applications. Having additional screen space “with a dedicated purpose, always accessible with a glance” (Grudin, 2001) reduces these burdens (Czerwinski et al., 2003), and studies suggest that providing multiple, distinct foci for interaction may aid users’ memory and recall (Tan, Stefanucci, Proffitt & Pausch, 2001; Tan, Stefanucci, Proffitt & Pausch, 2002). Finally, small displays can be used in conjunction with larger displays (Myers, Stiel & Gargiulo, 1998; Myers, 2000; Rekimoto, 1998), with controls and private information on the small device, and shared public information on the larger display. This shows how displays of different dimensions support completely different user activities and social conventions. It is also possible to dynamically join multiple displays for collaboration or to create a larger but temporary tiled display (Tandler, Prante, Müller-Tomfelde, Streitz & Steinmetz, 2001; Hinckley, 2003b; Hinckley, 2003a).

Large-Format Displays

Trends in display technology suggest that large-format displays will become increasingly affordable and common. A recent journal special issue includes numerous articles on implementing large-format displays using projection, application design for large displays, and specific application domains such as automotive design (Funkhouser & Li, 2000). Large displays often implicitly suggest multiple simultaneous users, with many applications revolving around collaboration (Swaminathan & Sato, 1997; Funkhouser & Li, 2000) and giving a large-scale

physical presence to virtual activities (Buxton, Fitzmaurice, Balakrishnan & Kurtenbach, 2000). To support input directly on whiteboard-size displays, researchers have explored gestural interaction techniques for pens or touchscreens (Guimbretiere, Stone & Winograd, 2001; Moran et al., 1997). Some technologies cannot handle more than one point of contact, so check this carefully if simultaneous use by multiple persons is desired. Large displays also seem to lend themselves to interaction at a distance, although using laser pointers to support such interaction (Olsen & Nielsen, 2001) has met with mixed success due to the lack of separation between tracking versus dragging states (Buxton, 1990b); using small handheld devices to interact with the full area of a large display also is problematic as the ratio of the display size to the control surface size may be very large (Myers et al., 1998). Environmentally situated ambient displays share some properties of large displays, but emphasize subtle presentation of information in the periphery of attention (Ishii & Ullmer, 1997; Wisneski et al., 1998). Large-format displays and virtual realities also share some design issues; see the taxonomy of Buxton & Fitzmaurice for further discussion (Buxton & Fitzmaurice, 1998).

Unless life-size viewing of large objects is necessary (Buxton et al., 2000), in general it is not yet clear what performance benefits a single large display may offer as compared to multiple monitors with the same screen area partitioned by bezels (Czerwinski et al., 2003). One recent study suggests that the increased field-of-view afforded by large-format displays can lead to improved 3D navigation performance, especially for women (Czerwinski, Tan & Robertson, 2002).

Force and Tactile Displays

Haptic feedback research has sought to provide an additional channel of sensory feedback by synthesizing forces on the skin of the operator. The touch sensation is extraordinarily complex. In fact, the sense of “touch” is a very imprecise term: it includes an amalgamation of multiple sensory systems, including sensitivity to pressure, small shear forces in the skin, heat and cold, pain, kinesthesia and proprioception, and the vestibular system (Mackenzie & Iberall, 1994; Burdea, 1996).

There appears to be no physical means by which a complex tactile stimulus can be delivered except in a very localized way. As a result, most haptic feedback devices are limited to simulation of a single point of contact, analogous to feeling the world with the tip of a pencil, although a few examples of whole-hand force feedback devices exist (Iwata, 1990; Burdea, 1996). Efforts in haptic feedback include force feedback (active presentation of forces to the user) and tactile feedback (active presentation of vibrotactile stimuli to the user). Haptic feedback is popular for gaming devices, such as force feedback steering wheels and joysticks, but general-purpose pointing devices with force or tactile feedback remain uncommon. For a comprehensive discussion of force and tactile feedback technologies and techniques, as well as perceptual properties of the skin and joints, see (Burdea, 1996).

Adding force feedback to a mouse or stylus may impose constraints on the mechanical design,

since a physical linkage is typically needed to reflect true forces. This may prevent a force feedback mouse from functioning like a traditional mouse, as it may limit range of motion or preclude clutching by lifting the device. Some devices instead increase resistance between the mouse and the pad, but this prevents simulation of hard contact forces. One can also use a vibrotactile stimulus, such as a vibrating pin under the mouse button, or vibrating the shaft of an isometric joystick (Campbell, Zhai, May & Maglio, 1999). Combination devices have also been explored (Akamatsu & Mackenzie, 1996). Vibrotactile feedback seems especially promising for small mobile devices, for example to provide the user with feedback of command recognition when the user's attention may not be focused on the screen (Poupyrev, Maruyama & Rekimoto, 2002). Applications for remote controls and augmented handles also look promising (Snibbe & MacLean, 2001; MacLean, Snibbe & Levin, 2000).

Using force feedback to provide attractive forces that pull the user towards a target, or tactile feedback to provide additional feedback for the boundaries of the target, has been found to yield modest speed improvements in some target acquisition experiments, although error rates may also increase (Akamatsu & Mackenzie, 1996; MacKenzie, 1995). However, there have been almost no published studies for tasks where multiple targets are present, as on a computer screen with many icons and menus. Haptic feedback for one target may interfere with the selection of another, unless one uses techniques such as reducing the haptic forces during rapid motion (Oakley, Brewster & Gray, 2001). Finally, one should also consider whether software constraints, such as snap-to grids, are sufficient to support the user's tasks.

The construction of force output devices is extremely technically demanding. They must be stiff in order to be able to create the sensation of solid contact, yet light so that they have little inertia themselves, and there must be a tight loop between input (position) and output (force). Sigoma (Sigoma, 1993) has suggested that having this loop iterate at 5 kHz may be necessary for optimal fine motor control. It has been shown that force feedback improves performance in certain telerobotic applications when, for example, inserting a peg into a hole (Sheridan, 1992). The most promising applications of force output seem to appear in domains where simulation of force is essential, such as surgical simulation and telerobotics (Burdea, 1996).

Another fundamental challenge for haptic feedback techniques results from the interaction between the haptic and visual channels. Visual dominance deals with phenomena resulting from the tendency for vision to dominate other modalities (Wickens, 1992). (Campbell et al., 1999) show that tactile feedback improves steering through a narrow tunnel, but only if the visual texture matches the tactile texture; otherwise tactile feedback harms performance.

Auditory Displays

Here, we consider computer-generated auditory feedback. Audio can consist of synthesized or recorded speech. All other audio feedback is known as nonspeech audio. With stereo speakers or a stereo headset, either type of audio can be presented such that it seems to come from a specific 3D location around the user, known as spatialized audio. For speech input and technology-

mediated human-human communication applications that treat stored voice as data, see the Voice and Speech section earlier in this chapter.

Nonspeech Audio

Nonspeech auditory feedback is prevalent in video games but largely absent from other interaction with computing devices. Providing an auditory echo of the visual interface has little or no practical utility, and may annoy users. Audio should be reserved to communicate simple, short messages that complement visual feedback (if any). Furthermore, one or more of the following conditions should hold: the message should (1) deal with events in time, (2) call for immediate action, or (3) take place when the user's visual attention may be overburdened or directed elsewhere (Deatherage, 1972; Buxton, 1995b). For example, researchers have attempted to enhance scrollbars using audio feedback (Brewster, Wright & Edwards, 1994). However, the meaning of such sounds may not be clear. Gaver advocates ecological sounds that resemble real-world events with an analogous meaning. For example, an empty disc drive might sound like a hollow metal container (Gaver, 1989). If a long or complex message must be delivered using audio, it will likely be quicker and clearer to deliver it using speech output. Audio feedback may be crucial to support tasks or functionality on mobile devices which must take place when the user is not looking at the display (for some examples, see (Hinckley et al., 2000)).

Nonspeech sounds can be especially useful for attracting the attention of the user. Auditory alerting cues have been shown to work well, but only in environments where there is low

auditory clutter. However, the number of simple nonspeech alerting signals is limited, and this can easily result in misidentification or cause signals to mask one another. An analysis of sound signals in fighter aircraft (Doll & Folds, 1985) found that the ground proximity warning and the angle-of-attack warning on an F16 were both an 800-Hz tone, a dangerous confound since these conditions require opposite responses from the pilot. It can also be difficult to devise nonspeech audio events that convey information without provoking an alerting response that unnecessarily interrupts the user. For example, this design tension arises when considering nonspeech audio cues that convey various properties of an incoming email message (Sawhney & Schmandt, 2000) (Hudson & Smith, 1996).

Speech Output

Speech auditory output is generally delivered either through recorded speech segments or completely synthetic speech (also known as text-to-speech technology). There has been considerable interest, especially for military applications, in the use of speech in providing warnings to the operators of complex systems. Speech can provide information to direct the operator's attention in a way that alarms cannot (since an unfamiliar alarm simply indicates a problem, without telling the user the nature or context of the problem). Synthetic speech is most useful where visual information is not available, for example, in touch-tone phone menu systems, or in screen reader software for blind or low-sighted users. Although progress is being made, synthetic voices still sound somewhat unnatural and may be more difficult for users to understand. Recorded speech is often used to give applications, particularly games, a more

personal feel, but can only be used for a limited number of responses known in advance.

The rate at which words must be produced to sound natural is a narrow range. For warning messages, 178 words per minute is intelligible but hurried, 123 words per minute is distracting and irritatingly slow, and a more natural rate of 156 words per minute is preferred (Simpson & Marchionda-Frost, 1984). The playback rate of speech can be increased by overlapping samples in time such that one sample is presented to one ear, and another sample to the other ear.

Technologies to correct for pitch distortions and remove pauses have also been developed (Arons, 1993; Stifelman, 1996; Sawhney & Schmandt, 2000). It is recommended by the U.S. Air Force that synthetic speech be 10 dB above ambient noise levels (Stokes et al., 1990).

Spatialized Audio Displays

It is possible to synthesize spatially localized sounds with a quality such that spatial localization in the virtual space is almost as good as localization of sounds in the natural environment (Wenzel, 1992). Auditory localization appears to be primarily a two-dimensional phenomenon. That is, observers can localize in horizontal position (azimuth) and elevation angle to some degree of accuracy. Azimuth and elevation accuracies are of the order of 15°. As a practical consequence this means that sound localization is of little use in identifying sources in conventional screen displays. Where localized sounds are really useful is in providing an orienting cue or warning about events occurring behind the user, outside of the field of vision. There is also a well-known phenomenon called visual capture of sound. Given a sound and an

apparent visual source for the sound, for example, a talking face on a cinema screen, the sound is perceived to come from the source despite the fact that the actual source may be off to one side. Thus, visual localization tends to dominate auditory localization when both kinds of cues are present.

Future Directions

The future of interaction with computers will both be very different and very much like it is today. Some of our current tools, such as mice and keyboards, have evolved to suit interaction with desktop graphical user interfaces and rapid text entry. As long as users' work continues to involve tasks such as calculating budgets, writing reports, looking up citations, exchanging memos, and other knowledge worker tasks that seem to lend themselves to solution using desktop computers, we will continue to see mice and keyboards in use, not only because they are familiar, but also because they closely match human skills and the requirements of the tasks. Devising new techniques that provide more efficient pointing at a desktop display than a mouse, for example, is difficult to achieve (Card et al., 1978). Speech recognition will allow new types of interaction and may enable interaction where it previously has been difficult or infeasible. However, even as technical limitations are removed, speech interaction will not replace all forms of interaction: we will continue to interact with computers using our hands and physical intermediaries, not necessarily because our technology requires us to do so, but because touching, holding, and moving physical objects is the foundation of the long evolution of tool use in the human species (Wilson, 1998).

But our computers and the tasks they serve are rapidly evolving. Current handheld devices have the display and computational capabilities of common desktop machines from several years ago. What is lacking is new methods of interacting with such devices that uniquely suit mobile interaction, rather than derivatives of the desktop interface. Researchers are still actively exploring and debating the best ways to achieve this. Meanwhile, technology advances and economic trends continue to drive the cost of commodity displays lower and lower, while the limits of the technology continue to increase. Thus we will continue to see new innovations in both very small and very large displays, and as these become commonplace new forms of interaction will become prevalent. Very small displays invariably seem to be incorporated into input/output appliances such as watches, pagers, and handheld devices, so interaction techniques for very small form-factors will become increasingly important.

The internet and wireless networking seem to be the main disruptive technologies of the current era. Indeed, it seems likely that 100 years from now the phrase “wireless network” will seem every bit as antiquated as the phrase “horseless carriage” does today. Nobody really understands yet what it will mean for everything and everyone to be connected, but many researchers are working to explore the vision of ubiquitous computing originally laid out by Mark Weiser (Weiser, 1991). Techniques that allow users to communicate and share information will become increasingly important. Biometric sensors or other convenient means for establishing identity will make services such as personalization of the interface and sharing data much simpler

(Rekimoto, 1997; Sugiura & Koseki, 1998). Techniques that combine dissimilar input devices and displays in interesting ways also will be important to realize the full potential of these technologies (e.g. (Myers, Miller, Evankovich & Bostwick, 2001; Streitz et al., 1999)).

Electronic tagging techniques for identifying objects (Want, Fishkin, Gujar & Harrison, 1999) may also become commonplace. Such a diversity of locations, users, and task contexts points to the increasing importance of sensors to acquire contextual information, as well as machine learning techniques to interpret them and infer meaningful actions (Buxton, 1995a; Bellotti et al., 2002; Hinckley et al., 2003). This may well lead to an age of ubiquitous sensors (Saffo, 1997) with devices that can see, feel, and hear through digital perceptual mechanisms.

Defining Terms

Absolute input device An input device that reports its actual position, rather than relative movement. A tablet or touchscreen typically operates this way (see also relative input device).

Antialiasing The specification of pixel color values so that they reflect the correct proportions of the colored regions that contribute to that pixel. In temporal antialiasing the amount of time a region of a simulated scene contributes to a pixel is also taken into account.

Acquisition time The average time to pick up or put down an input device. Sometimes known as homing time.

Augmented reality The superimposition of artificially generated graphical elements on objects in the environment. Achieved with a see-through head mounted display.

Background sensing techniques Implicitly sensed interaction takes place in the background, behind the fore of the user's attention. Background sensing techniques use sensor technology or intelligent algorithms to glean additional, typically neglected, information from the existing input stream, with the goal of supporting the user with semi-automatic or implicit actions and services.

Cognitive chunk. A series of elemental tasks that seems like a single concept to the user. For example, users think of pointing at something as a single chunk, but from a technical perspective it may consist of selecting an (X, Y, Z) coordinate in a 3D environment. By using technologies and interaction metaphors that parallel the way the user thinks about a task as closely as possible, the designer can phrase together a series of elemental tasks into a single cognitive chunk.

Compound tasks A compound task is a hierarchy of elemental sub-tasks. For example, the navigate/select compound task consists of scrolling to view an item in a list, and then clicking on it to select it. When interacting with a graphical scroll bar, scrolling itself may be a compound task with multiple selection or positioning tasks.

Control-to-display (C:D) ratio The ratio between the movement a user must make with an input device and the resulting movement obtained on the display. With a large C:D ratio, a large

movement is required to effect a small change on the display, affording greater precision. A low ratio allows more rapid operation and takes less desk space. The C:D ratio is sometimes expressed as a single number, in which case it is referred to as the device gain. Note that many experts have criticized gain as a fundamental concept; one must take great care when manipulating gain in experiments, since it confounds display size and control size in one arbitrary metric.

Direct input device A device that the user operates directly on the screen or other display to be controlled, such as a touch screen (see also indirect input device).

Fish tank virtual reality A form of virtual reality display that confines the virtual scene to the vicinity of a monitor screen.

Fitts' Law A model that relates the movement time to point at a target, the amplitude of the movement (the distance to the target), and the width of the target (i.e., the precision requirement of the pointing movement). The movement time is proportional to the logarithm of the distance divided by the target width, with constant terms that vary from one device to another. Fitts' Law has found wide application in HCI to evaluating and comparing input devices and transfer functions for pointing at targets.

Flicker fusion frequency The frequency at which a flickering light is perceived as a steady

illumination. Useful in determining the requirements for a visual display.

Footprint The physical movement space (area) required to operate an input device.

Fovea The central part of the retina at which vision is the sharpest. About 2° of visual angle in diameter.

Gamma correction The correction of nonlinearities of a monitor so that it is possible to specify a color in linear coordinates.

Indirect input device A device that the user operates by moving a control that is located away from the screen or other display to be controlled, such as a mouse or trackball (see also direct input device).

Input device A hardware computer peripheral through which the user interacts with the computer.

Interaction task A low-level primitive input to be obtained from the user, such as entering a text string or choosing a command.

Interaction technique The fusion of input and output, consisting of all hardware and software

elements, that provides a particular way for the user to accomplish a low-level task with a physical input device. For example, the pop-up menu is an interaction technique for choosing a command or other item from a small set, using a mouse and a graphical display.

Lambertian diffuser A diffuser that spreads incoming light equally in all directions.

Latency The end-to-end delay between the user's physical movement, and the system's ultimate feedback to the user. Latency of more than 75-100 milliseconds significantly impairs user performance for many interactive tasks.

Luminance The standard way of defining an amount of light. This measure takes into account the relative sensitivities of the human eye to light of different wavelengths.

Preattentive processing Visual stimuli that are processed at an early stage in the visual system in parallel. This processing is done prior to processing by the mechanisms of visual attention.

Refresh rate The rate at which a computer monitor is redrawn. Sometimes different from the update rate.

Relative input device An input device that reports its distance and direction of movement each time it is moved, but cannot report its absolute position. A mouse operates this way (see absolute

input device).

Screen gain A measure of the amount by which a projection video screen reflects light in a preferred direction. The purpose is to give brighter images if viewed from certain positions. There is a corresponding loss in brightness from other viewing positions.

Superacuties The ability to perceive visual effects with a resolution that is finer than can be predicted from the spacing of receptors in the human eye.

Transfer function A mathematical transformation that scales the data from an input device to ideally provide smooth, efficient, and intuitive operation. Appropriate mappings are transfer functions that match the physical properties sensed by the input device, and include force-to-velocity, position-to-position, and velocity-to-velocity functions.

Three-state model A model for the discrete states of input devices which models transitions between three states: tracking, dragging, and out-of-range. Most input devices only sense two of these three states (for example, a mouse senses tracking and dragging, whereas a touchpad senses tracking and the out-of-range state).

Uniform color space A transformation of a color specification such that equal metric differences between colors more closely correspond to equal perceptual differences.

Update rate The rate at which the image on a computer monitor is changed.

Virtual reality A method of monitoring a users head position and creating a perceptive view of an artificial world that changes as the user moves, in such a way as to simulate an illusory three-dimensional scene.

Visual acuity The ability of the human visual system to resolve fine targets.

==

References

- Accot, J. & S. Zhai (1997). Beyond Fitts' Law: Models for Trajectory-Based HCI Tasks. Proc. CHI'97: ACM Conference on Human Factors in Computing Systems. 295-302.
- Accot, J. & S. Zhai (1999). Performance Evaluation of Input Devices in Trajectory-based Tasks: An Application of the Steering Law. Proc. CHI'99, Pittsburgh, PA. 466-472.
- Accot, J. & S. Zhai (2001). Scale Effects in Steering Law Tasks. Proc. CHI'2001 ACM Conference on Human Factors in Computing Systems. 1-8.
- Accot, J. & S. Zhai (2002). More than dotting the i's-- Foundations for crossing-based interfaces. ACM CHI 2002 Conf. on Human Factors in Computing Systems. 73-80.
- Ahlberg, C. & B. Shneiderman (1994). The alphaslider: a compact and rapid selector. CHI'94. 365-371.
- Akamatsu, M. & I. S. Mackenzie (1996). "Movement Characteristics Using a Mouse with Tactile and Force Feedback." International Journal of Human-Computer Studies **45**: 483-493.
- Anderson, J. R. (1980). Chapter 8: Cognitive Skills. Cognitive Psychology and Its Implications. San Francisco, W. H. Freeman: 222-254.
- Arons, B. (1993). SpeechSkimmer: Interactively Skimming Recorded Speech. UIST'93 Symp. on User Interface Software & Technology. 187-195.
- Balakrishnan, R., T. Baudel, G. Kurtenbach & G. Fitzmaurice (1997). The Rockin'Mouse: Integral 3D Manipulation on a Plane. CHI'97 Conf. on Human Factors in Computing Systems. 311-318.
- Balakrishnan, R. & K. Hinckley (1999). The Role of Kinesthetic Reference Frames in Two-Handed Input Performance. Proc. ACM UIST'99 Symp. on User Interface Software and Technology. 171-178.
- Balakrishnan, R. & K. Hinckley (2000). Symmetric Bimanual Interaction. CHI 2000. 33-40.
- Balakrishnan, R. & G. Kurtenbach (1999). Exploring Bimanual Camera Control and Object Manipulation in 3D Graphics Interfaces. Proc. CHI'99 ACM Conf. on Human Factors in Computing Systems. 56-63.
- Balakrishnan, R. & I. S. MacKenzie (1997). Performance Differences in the Fingers, Wrist, and Forearm in Computer Input Control. Proc. CHI'97 ACM Conf. on Human Factors in Computing Systems. 303-310.
- Baudel, T. & M. Beaudouin-Lafon (1993). "Charade: Remote Control of Objects Using Hand Gestures." Communications of the ACM **36**(7): 28-35.
- Baudisch, P., N. Good, V. Bellotti & P. Schraedley (2002). Keeping Things in Context: A Comparative Evaluation of Focus Plus Context Screens, Overviews, and Zooming. CHI 2002 Conf. on Human Factors in Computing Systems. 259-266.

- Bederson, B. (2000). Fisheye Menus. UIST 2000. 217-226.
- Bederson, B., J. Hollan, K. Perlin, J. Meyer, D. Bacon & G. Furnas (1996). "Pad++: A Zoomable Graphical Sketchpad for Exploring Alternate Interface Physics." Journal of Visual Languages and Computing 7: 3-31.
- Bellotti, V., M. Back, W. K. Edwards, R. Grinter, C. Lopes & A. Henderson (2002). Making sense of sensing systems: Five questions for designers and researchers. Proc. ACM CHI 2002 Conference on Human Factors in Computing Systems. 415-422.
- Betrissey, C., J. Blinn, B. Dresevic, B. Hill, G. Hitchcock, B. Keely, D. Mitchell, J. Platt & T. Whitted (2000). Displaced Filtering for Patterned Displays. Proc. Society for Information Display Symposium. 296-299.
- Bier, E., M. Stone, K. Pier, W. Buxton & T. DeRose (1993). Toolglass and Magic Lenses: The See-Through Interface. Proceedings of SIGGRAPH 93, Anaheim, Calif. 73-80.
- Bolt, R. (1980). "Put-That-There: Voice and Gesture at the Graphics Interface." Computer Graphics(Aug.): 262-270.
- Brewster, S. A., P. C. Wright & A. D. N. Edwards (1994). The design and evaluation of an auditory-enhanced scrollbar. Conference proceedings on Human factors in computing systems. 173 -179.
- Britton, E., J. Lipscomb & M. Pique (1978). "Making Nested Rotations Convenient for the User." Computer Graphics 12(3): 222-227.
- Brooks, J., F. P. (1988). Grasping reality through illusion: interactive graphics serving science. Proceedings of CHI'88: ACM Conference on Human Factors in Computing Systems, Washington, DC, ACM, New York. 1-11.
- Bukowski, R. & C. Sequin (1995). Object Associations: A Simple and Practical Approach to Virtual 3D Manipulation. ACM 1995 Symposium on Interactive 3D Graphics. 131-138.
- Burdea, G. (1996). Force and Touch Feedback for Virtual Reality. New York, NY, John Wiley & Sons.
- Buxton, B. & G. Fitzmaurice (1998). "HMDs, Caves & Chameleon: A Human-Centric Analysis of Interaction in Virtual Space." Computer Graphics 32(8): 69-74.
- Buxton, W. (1983). "Lexical and Pragmatic Considerations of Input Structure." Computer Graphics 17(1): 31-37.
- Buxton, W. (1986). Chunking and Phrasing and the Design of Human-Computer Dialogues. Information Processing '86, Proc. of the IFIP 10th World Computer Congress, Amsterdam: North Holland Publishers. 475-480.
- Buxton, W. (1990a). The Pragmatics of Haptic Input. Proceedings of CHI'90: ACM Conference on Human Factors in Computing Systems, Tutorial 26 Notes, Seattle, Wash., ACM, New York.
- Buxton, W. (1990b). A three-state model of graphical input. Proc. INTERACT'90, Amsterdam: Elsevier Science. 449-456.

- Buxton, W. (1995a). Integrating the Periphery and Context: A New Taxonomy of Telematics. Proceedings of Graphics Interface '95. 239-246.
- Buxton, W. (1995b). Speech, Language and Audition. Readings in Human-Computer Interaction: Toward the Year 2000. R. Baecker, J. Grudin, W. Buxton and S. Greenberg, Morgan Kaufmann Publishers: 525-537.
- Buxton, W., G. Fitzmaurice, R. Balakrishnan & G. Kurtenbach (2000). "Large displays in automotive design." IEEE Computer Graphics and Applications(July/August): 68-75.
- Buxton, W., E. Fiume, R. Hill, A. Lee & C. Woo (1983). Continuous hand-gesture driven input. Proceedings of Graphics Interface '83. 191-195.
- Buxton, W., R. Hill & P. Rowley (1985). "Issues and Techniques in Touch-Sensitive Tablet Input." Computer Graphics **19**(3): 215-224.
- Buxton, W. & B. Myers (1986). A Study in Two-Handed Input. Proceedings of CHI'86: ACM Conference on Human Factors in Computing Systems, Boston, Mass., ACM, New York. 321-326.
- Buyukkokten, O., H. Garcia-Molina & A. Paepcke (2001). Accordion Summarization for End-Game Browsing on PDAs and Cellular Phones. ACM CHI 2001 Conf. on Human Factors in Computing Systems, Seattle, WA.
- Cadoz, C. (1994). Les realites virtuelles, Dominos, Flammarion.
- Campbell, C., S. Zhai, K. May & P. Maglio (1999). What You Feel Must Be What You See: Adding Tactile Feedback to the Trackpoint. Proceedings of INTERACT'99: 7th IFIP conference on Human Computer Interaction. 383-390.
- Card, S., W. English & B. Burr (1978). "Evaluation of mouse, rate-controlled isometric joystick, step keys, and text keys for text selection on a CRT." Ergonomics **21**: 601-613.
- Card, S., J. Mackinlay & G. Robertson (1991). "A Morphological Analysis of the Design Space of Input Devices." ACM Transactions on Information Systems **9**(2): 99-122.
- Card, S., J. Mackinlay & B. Shneiderman (1999). Readings in Information Visualization: Using Vision to Think. San Francisco, Morgan Kaufmann.
- Card, S., T. Moran & A. Newell (1980). "The Keystroke-Level Model for User Performance Time with Interactive Systems." Communications of the ACM **23**(7): 396-410.
- Cassell, J. (2003). A Framework for Gesture Generation and Interpretation. Computer Vision in Human-Machine Interaction. R. Cipolla and A. Pentland, Cambridge University Press: (in press).
- Caudell, T. P. & D. W. Mizell (1992). Augmented Reality: An application of heads-up display technology to manual manufacturing processes. Proc. HICCS '92.

- Chance, S., F. Gaunet, A. Beall & J. Loomis (1998). "Locomotion Mode Affects the Updating of Objects Encountered During Travel: The Contribution of Vestibular and Proprioceptive Inputs to Path Integration." Presence 7(2): 168-178.
- Chen, M., S. J. Mountford & A. Sellen (1988). "A Study in Interactive 3-D Rotation Using 2-D Control Devices." Computer Graphics 22(4): 121-129.
- Cholewiak, R. & A. Collins (1991). Sensory and physiological bases of touch. The psychology of touch. M. Heller and W. Schiff, Kawrence Erlbaum: 23-60.
- Christ, R. E. (1975). "Review and analysis of color coding research for visual displays." Human Factors 25: 71-84.
- Cohen, P., M. Johnston, D. McGee, S. Oviatt, J. Pittman, I. Smith, L. Chen & J. Clow (1997). QuickSet: Multimodal Interaction for Distributed Applications. ACM Multimedial 97.
- Cohen, P. R. & J. W. Sullivan (1989). Synergistic Use of Direct Manipulation and Natural Language. Proc. ACM CHI'89 Conference on Human Factors in Computing Systems. 227-233.
- Cole, W. G. (1986). Medical cognitive graphics. ACM CHI'86 Conf. on Human factors in Computing Systems. 91-95.
- Conner, D., S. Snibbe, K. Herndon, D. Robbins, R. Zeleznik & A. van Dam (1992). Three-Dimensional Widgets. Computer Graphics (Proc. 1992 Symposium on Interactive 3D Graphics). 183-188, 230-231.
- Cook, R. L. (1986). "Stochastic sampling in computer graphics." ACM Trans. Graphics 5(1): 51-72.
- Cowan, W. B. (1983). "An inexpensive calibration scheme for calibrations of a color monitor in terms of CIE standard coordinates." Computer Graphics 17(3): 315-321.
- Czerwinski, M., G. Smith, T. Regan, B. Meyers, G. Robertson & G. Starkweather (2003). Toward Characterizing the Productivity Benefits of Very Large Displays. To appear in INTERACT 2003.
- Czerwinski, M., D. S. Tan & G. G. Robertson (2002). Women Take a Wider View. Proc. ACM CHI 2002 Conference on Human Factors in Computing Systems, Minneapolis, MN. 195-202.
- Darken, R. P. & J. L. Sibert (1993). A Toolset for Navigation in Virtual Environments. UIST'93. 157-165.
- Darken, R. P. & J. L. Sibert (1995). "Navigating Large Virtual Spaces." International Journal of Human-Computer Interaction(Oct.).
- Deatherage, B. H. (1972). Auditory and Other Sensory Forms of Information Presentation. Human Engineering Guide to Equipment Design. H. Van Cott and R. Kinkade, U.S. Government Printing Office.
- Deering, M. (1992). "High resolution virtual reality." Computer Graphics 26(2): 195-202.

- Dey, A., G. Abowd & D. Salber (2001). "A Conceptual Framework and a Toolkit for Supporting the Rapid Prototyping of Context-Aware Applications." Journal of Human-Computer Interaction **16**(2-4): 97-166.
- Doll, T. J. & D. J. Folds (1985). Auditory signals in military aircraft: ergonomic principles versus practice. Proc. 3rd Symp. Aviation Psych., Ohio State University, Dept. of Aviation, Columbus, OH. 111-125.
- Douglas, S., A. Kirkpatrick & I. S. MacKenzie (1999). Testing Pointing Device Performance and User Assessment with the ISO 9241, Part 9 Standard. Proc. ACM CHI'99 Conf. on Human Factors in Computing Systems. 215-222.
- Douglas, S. & A. Mithal (1994). The Effect of Reducing Homing Time on the Speed of a Finger-Controlled Isometric Pointing Device. Proc. ACM CHI'94 Conf. on Human Factors in Computing Systems. 411-416.
- Feiner, S., B. Macintyre & D. Seligmann (1993). "Knowledge-Based Augmented Reality." Communications of the ACM **36**(7): 53-61.
- Fitts, P. (1954). "The information capacity of the human motor system in controlling the amplitude of movement." Journal of Experimental Psychology **47**: 381-391.
- Fitzmaurice, G. & W. Buxton (1997). An Empirical Evaluation of Graspable User Interfaces: towards specialized, space-multiplexed input. Proceedings of CHI'97: ACM Conference on Human Factors in Computing Systems, Atlanta, Georgia, ACM, New York. 43-50.
- Fitzmaurice, G., H. Ishii & W. Buxton (1995). Bricks: Laying the Foundations for Graspable User Interfaces. Proceedings of CHI'95: ACM Conference on Human Factors in Computing Systems, Denver, Colorado, ACM, New York. 442-449.
- Fitzmaurice, G. W., R. Balakrishnan & G. Kurtenbach (1999). "Sampling, synthesis, and input devices." Commun. ACM **42**(8): 54 - 63.
- Foley, J. D., V. Wallace & P. Chan (1984). "The Human Factors of Computer Graphics Interaction Techniques." IEEE Computer Graphics and Applications(Nov.): 13-48.
- Freeman, W. T. & C. Weissman (1995). Television control by hand gestures. Intl. Workshop on Automatic Face and Gesture Recognition, Zurich, Switzerland. 179-183.
- Funkhouser, T. & K. Li (2000). "Large Format Displays." IEEE Comput. Graphics Appl.(July-Aug. special issue): 20-75.
- Gaver, W. (1989). "The SonicFinder: An interface that uses auditory icons." Human-Computer Interaction **4**(1): 67-94.
- Gibson, J. (1986). The Ecological Approach to Visual Perception, Lawrence Erlbaum Assoc. Hillsdale, NJ.
- Goldberg, D. & C. Richardson (1993). Touch-Typing with a Stylus. Proc. INTERCHI'93 Conf. on Human Factors in Computing Systems. 80-87.

- Green, M. & J. Liang (1994). "JDCAD: A highly interactive 3D modeling system." Computers and Graphics **18**(4): 499-506.
- Grudin, J. (2001). Partitioning Digital Worlds: Focal and Peripheral Awareness in Multiple Monitor Use. CHI 2001. 458-465.
- Guiard, Y. (1987). "Asymmetric Division of Labor in Human Skilled Bimanual Action: The Kinematic Chain as a Model." The Journal of Motor Behavior **19**(4): 486-517.
- Guiard, Y., F. Buourgeois, D. Mottet & M. Beaudouin-Lafon (2001). Beyond the 10-bit Barrier: Fitts' Law in Multi-Scale Electronic Worlds. IHM-HCI 2001, Lille, France.
- Guimbretiere, F., M. C. Stone & T. Winograd (2001). Fluid Interaction with High-resolution Wall-size Displays. Proc. UIST 2001 Symp. on User Interface Software and Technology. 21-30.
- Harrison, B., K. Fishkin, A. Gujar, C. Mochon & R. Want (1998). Squeeze Me, Hold Me, Tilt Me! An Exploration of Manipulative User Interfaces. Proc. ACM CHI'98 Conf. on Human Factors in Computing Systems. 17-24.
- Harrison, B., H. Ishii, K. Vicente & W. Buxton (1995). Transparent Layered User Interfaces: An Evaluation of a Display Design to Enhance Focused and Divided Attention. Proceedings of CHI'95: ACM Conference on Human Factors in Computing Systems. 317-324.
- Harrison, B., G. Kurtenbach & K. Vicente (1995). An Experimental Evaluation of Transparent User Interface Tools and Information Content. UIST'95. 81-90.
- Harrison, B. & K. Vicente (1996). An Experimental Evaluation of Transparent Menu Usage. Proceedings of CHI'96: ACM Conference on Human Factors in Computing Systems. 391-398.
- Hauptmann, A. (1989). Speech and Gestures for Graphic Image Manipulation. Proceedings of CHI'89: ACM Conference on Human Factors in Computing Systems, Austin, Texas, ACM, New York. 241-245.
- Hinckley, K. (2003a). Distributed Sensing Techniques for Face-to-Face Collaboration. ICMI-PUT'03 Fifth International Conference on Multimodal Interfaces, Vancouver B.C., Canada.
- Hinckley, K. (2003b). Synchronous Gestures for Multiple Users and Computers. currently submitted to UIST'03 Symp. On User Interface Software & Technology. 10 pp. (full paper).
- Hinckley, K., E. Cutrell, S. Bathiche & T. Muss (2001). Quantitative Analysis of Scrolling Techniques. currently submitted to CHI 2002.
- Hinckley, K., M. Czerwinski & M. Sinclair (1998a). Interaction and Modeling Techniques for Desktop Two-Handed Input. Proceedings of the ACM UIST'98 Symposium on User Interface Software and Technology, San Francisco, Calif., ACM, New York. 49-58.

- Hinckley, K., R. Pausch, J. Goble & N. Kassell (1994a). Passive real-world interface props for neurosurgical visualization. Proceedings of CHI'94: ACM Conference on Human Factors in Computing Systems, Boston, Mass., ACM, New York. 452-458.
- Hinckley, K., R. Pausch, J. C. Goble & N. F. Kassell (1994b). A Survey of Design Issues in Spatial Input. Proceedings of the ACM UIST'94 Symposium on User Interface Software and Technology, Marina del Rey, Calif., ACM, New York. 213-222.
- Hinckley, K., R. Pausch, D. Proffitt & N. Kassell (1998b). "Two-Handed Virtual Manipulation." ACM Transactions on Computer-Human Interaction 5(3): 260-302.
- Hinckley, K., J. Pierce, E. Horvitz & M. Sinclair (2003). "Foreground and Background Interaction with Sensor-Enhanced Mobile Devices." ACM TOCHI (submitted for review)(Special Issue on Sensor-Based Interaction).
- Hinckley, K., J. Pierce, M. Sinclair & E. Horvitz (2000). Sensing Techniques for Mobile Interaction. ACM UIST 2000 Symp. on User Interface Software & Technology. 91-100.
- Hinckley, K. & M. Sinclair (1999). Touch-Sensing Input Devices. ACM CHI'99 Conf. on Human Factors in Computing Systems. 223-230.
- Hinckley, K., M. Sinclair, E. Hanson, R. Szeliski & M. Conway (1999). The VideoMouse: A Camera-Based Multi-Degree-of-Freedom Input Device. ACM UIST'99 Symp. on User Interface Software & Technology. 103-112.
- Hinckley, K., J. Tullio, R. Pausch, D. Proffitt & N. Kassell (1997). Usability Analysis of 3D Rotation Techniques. Proc. ACM UIST'97 Symp. on User Interface Software and Technology, Banff, Alberta, Canada, ACM, New York. 1-10.
- Hix, D., J. Templeman & R. Jacob (1995). Pre-Screen Projection: From Concept to Testing of a New Interaction Technique. CHI'95. 226-233.
- Honan, M., E. Serina, R. Tal & D. Rempel (1995). Wrist Postures While Typing on a Standard and Split Keyboard. Proc. HFES Human Factors and Ergonomics Society 39th Annual Meeting. 366-368.
- Horvitz, E., J. Breese, D. Heckerman, D. Hovel & K. Rommelse (1998). The Lumiere Project: Bayesian User Modeling for Inferring the Goals and Needs of Software Users. Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence, Madison, WI, July 1998, pages 256-265. Morgan Kaufmann: San Francisco.
- Horvitz, E., A. Jacobs & D. Hovel (1999). Attention-Sensitive Alerting. Proceedings of UAI '99, Conference on Uncertainty and Artificial Intelligence. 305-313.
- Hudson, S. & I. Smith (1996). Electronic Mail Previews Using Non-Speech Audio. CHI'96 Companion Proceedings. 237-238.

- Igarashi, T., S. Matsuoka & H. Tanaka (1999). Teddy: A Sketching Interface for 3D Freeform Design. ACM SIGGRAPH'99, Los Angeles, CA. 409-416.
- Ishii, H. & B. Ullmer (1997). Tangible Bits: Towards Seamless Interfaces between People, Bits, and Atoms. Proceedings of CHI'97: ACM Conference on Human Factors in Computing Systems, Atlanta, Georgia, ACM, New York. 234-241.
- Iwata, H. (1990). "Artificial Reality with Force-feedback: Development of Desktop Virtual Space with Compact Master Manipulator." Computer Graphics **24**(4): 165-170.
- Jacob, R. (1991). "The Use of Eye Movements in Human-Computer Interaction Techniques: What You Look At is What You Get." ACM Transactions on Information Systems **9**(3): 152-169.
- Jacob, R., L. Sibert, D. McFarlane & M. Mullen, Jr. (1994). "Integrality and Separability of Input Devices." ACM Transactions on Computer-Human Interaction **1**(1): 3-26.
- Jellinek, H. & S. Card (1990). Powermice and User Performance. Proc. ACM CHI'90 Conf. on Human Factors in Computing Systems. 213-220.
- Jojic, N., B. Brumitt, B. Meyers & S. Harris (2000). Detecting and estimating pointing gestures in dense disparity maps. Proceed. of IEEE Intl. Conf. on Automatic Face and Gesture Recognition.
- Jones, M., G. Marsden, N. Mohd-Nasir, K. Boone & G. Buchanan (1999). "Improving Web interaction on small displays." Computer Networks **31**(11-16): 1129-1137.
- Kabbash, P., W. Buxton & A. Sellen (1994). Two-handed input in a compound task. Proceedings of CHI'94: ACM Conference on Human Factors in Computing Systems, Boston, Mass., ACM, New York. 417-423.
- Kamba, T., S. A. Elson, T. Harpold, T. Stamper & P. Sukaviriya (1996). Using small screen space more efficiently. Conference proceedings on Human factors in computing systems. 383.
- Karat, C., C. Halverson, D. Horn & J. Karat (1999). Patterns of Entry and Correction in Large Vocabulary Continuous Speech Recognition Systems. Proc. ACM CHI'99 Conf. on Human Factors in Computing Systems. 568-575.
- Kirsh, D. (1995). Complementary strategies: why we use our hands when we think. Proceedings of 7th Annual Conference of the Cognitive Science Society., Hillsdale, NJ: Lawrence Erlbaum. 212-217.
- Kirsh, D. & P. Maglio (1994). "On Distinguishing Epistemic from Pragmatic Action." Cognitive Science **18**(4): 513-549.
- Kramer, A. (1994). Translucent Patches--Dissolving Windows. Proc. ACM UIST'94 Symp. on User Interface Software & Technology. 121-130.
- Kurtenbach, G. & W. Buxton (1991). Issues in Combining Marking and Direct Manipulation Techniques. Proc. UIST'91. 137-144.

- Kurtenbach, G. & W. Buxton (1993). The Limits of Expert Performance Using Hierarchic Marking Menus. Proc. INTERCHI'93. 482-487.
- Kurtenbach, G., G. Fitzmaurice, T. Baudel & B. Buxton (1997). The Design of a GUI Paradigm based on Tablets, Two-hands, and Transparency. Proceedings of CHI'97: ACM Conference on Human Factors in Computing Systems, Atlanta, Georgia, ACM, New York. 35-42.
- Kurtenbach, G., A. Sellen & W. Buxton (1993). "An empirical evaluation of some articulatory and cognitive aspects of 'marking menus'." Journal of Human Computer Interaction **8**(1).
- Lewis, J., K. Potosnak & R. Magyar (1997). Keys and Keyboards. Handbook of Human-Computer Interaction. M. Helander, T. Landauer and P. Prabhu. Amsterdam, North-Holland: 1285-1316.
- Lipscomb, J. & M. Pique (1993). "Analog Input Device Physical Characteristics." SIGCHI Bulletin **25**(3): 40-45.
- Loomis, J., R. L. Klatzky, R. G. Golledge & J. W. Philbeck (1999). Human navigation by path integration. Wayfinding: Cognitive mapping and other spatial processes. R. G. Golledge. Baltimore, Johns Hopkins: 125-151.
- Lucente, M., G. Zwart & A. George (1998). Visualization Space: A Testbed for Deviceless Multimodal User Interface. AAAI '98.
- Mackenzie, C. & T. Iberall (1994). The Grasping Hand. Amsterdam, North Holland.
- MacKenzie, I. S. (1992). "Fitts' law as a research and design tool in human-computer interaction." Human-Computer Interaction **7**: 91-139.
- MacKenzie, I. S. (1995). Input Devices and Interaction Techniques for Advanced Computing. Virtual environments and advanced interface design. W. Barfield and T. Furness. Oxford, UK, Oxford University Press: 437-470.
- MacKenzie, I. S. & Y. Guiard (2001). The Two-Handed Desktop Interface: Are We There Yet? Proc. ACM CHI 2001 Conf. on Human Factors in Computing Systems: Extended Abstracts. 351-352.
- MacKenzie, I. S., T. Kauppinen & M. Silfverberg (2001). Accuracy measures for evaluating computer pointing devices. CHI 2001. 9-16.
- MacKenzie, I. S. & A. Oniszczak (1998). A Comparison of Three Selection Techniques for Touchpads. Proc. ACM CHI'98 Conf. on Human Factors in Computing Systems. 336-343.
- MacKenzie, I. S., A. Sellen & W. Buxton (1991). A Comparison of Input Devices in Elemental Pointing and Dragging Tasks. Proc. ACM CHI '91 Conf. on Human Factors in Computing Systems. 161-166.
- MacKenzie, I. S. & R. W. Soukoreff (2002). A model of two-thumb text entry. Proceedings of Graphics Interface, Toronto, Canadian Information Processing Society. 117-124.

- MacKenzie, I. S. & C. Ware (1993). Lag as a Determinant of Human Performance in Interactive Systems. Proc. ACM INTERCHI'93 Conference on Human Factors in Computing Systems. 488-493.
- MacKenzie, I. S. & S. Zhang (1997). The Immediate Usability of Graffiti. Proc. Graphics Interface '97. 129-137.
- MacLean, K. E., S. S. Snibbe & G. Levin (2000). Tagged Handles: Merging Discrete and Continuous Control. Proc. ACM CHI 2000 Conference on Human Factors in Computing Systems, The Hague, Netherlands.
- Maes, P., T. Darrell, B. Blumberg & A. Pentland (1996). "The ALIVE system: wireless, full-body interaction with autonomous agents." ACM Multimedia Systems(Special Issue on Multimedia and Multisensory Virtual Worlds).
- Marklin, R. & G. Simoneau (1996). Upper extremity posture of typists using alternative keyboards. ErgoCon'96. 126-132.
- Marklin, R., G. Simoneau & J. Monroe (1997). The Effect of Split and Vertically-Inclined Computer Keyboards on Wrist and Forearm Posture. Proc. HFES Human Factors and Ergonomics Society 41st Annual Meeting. 642-646.
- Mathias, E., I. S. MacKenzie & W. Buxton (1996). "One-handed touch typing on a qwerty keyboard." Human Computer Interaction **11**(1): 1-27.
- McGuffin, M. & R. Balakrishnan (2002). "Acquisition of Expanding Targets." CHI Letters **4**(1).
- McLoone, H., K. Hinckley & E. Cutrell (2003). Bimanual Interaction on the Microsoft Office Keyboard. To appear in INTERACT 2003.
- Mine, M., F. Brooks & C. Sequin (1997). "Moving Objects in Space: Exploiting Proprioception in Virtual-Environment Interaction." Computer Graphics **31**(Proc. SIGGRAPH'97): 19-26.
- Moran, T., P. Chiu & W. van Melle (1997). Pen-Based Interaction Techniques for Organizing Material on an Electronic Whiteboard. Proc. ACM UIST'97 Symp. on User Interface Software & Technology. 45-54.
- Myers, B., Lie, K., Yang, B. (2000). Two-Handed Input using a PDA and a Mouse. CHI 2000. 41-48.
- Myers, B., R. Miller, C. Evankovich & B. Bostwick (2001). Individual Use of Hand-Held and Desktop Computers Simultaneously. submitted?
- Myers, B., H. Stiel & R. Gargiulo (1998). Collaboration Using Multiple PDAs Connected to a PC. Proc. ACM CSCW'98 Conf. on Computer Supported Cooperative Work, Seattle, WA. 285-294.
- Mynatt, E. D., T. Igarashi, W. K. Edwards & A. LaMarca (1999). Flatland: New Dimensions in Office Whiteboards. ACM SIGCHI Conference on Human Factors in Computing Systems, Pittsburgh, PA. 346-353.
- Nguyen, D. H. & E. Mynatt (2001). "Towards Visibility of a Ubicomp Environment."
- Norman, D. (1990). The Design of Everyday Things, Doubleday: New York, NY.

- Noyes, J. (1983). "Chord Keyboards." Applied Ergonomics **14**: 55-59.
- Oakley, I., S. Brewster & P. Gray (2001). Solving Multi-Target Haptic Problems in Menu Interaction. Proc. ACM CHI 2001 Conf. on Human Factors in Computing Systems: Extended Abstracts. 357-358.
- Olsen, D. R. & T. Nielsen (2001). Laser pointer interaction. Proc. ACM CHI 2001 Conf. on Human Factors in Computing Systems. 17-22.
- Olson, J. R. & G. M. Olson (1990). "The Growth of Cognitive Modeling in Human-Computer Interaction Since GOMS." Human-Computer Interaction **5**(2 and 3): 221-266.
- Oviatt, S. (1997). "Multimodal Interactive Maps: Designing for Human Performance." Human-Computer Interaction **12**: 93-129.
- Pearson, G. & M. Weiser (1988). Exploratory Evaluation of a Planar Foot-Operated Cursor-Positioning Device. Proc. ACM CHI'88 Conference on Human Factors in Computing Systems. 13-18.
- Pekelney, R. & R. Chu (1995). Design Criteria of an Ergonomic Mouse Computer Input Device. Proc. HFES Human Factors and Ergonomics Society 39th Annual Meeting. 369-373.
- Perlin, K. & D. Fox (1993). Pad: An Alternative Approach to the Computer Interface. SIGGRAPH '93.
- Platt, J. (2000). "Optimal Filtering for Patterned Displays." IEEE Signal Processing Letters **7**(7): 179-83.
- Plumlee, M. & C. Ware (2002). Modeling performance for zooming vs multi-window interfaces based on visual working memory. AVI '02: Advanced Visual Interfaces, Trento Italy.
- Poupyrev, I., S. Maruyama & J. Rekimoto (2002). Ambient touch: designing tactile interfaces for handheld devices. UIST 2002 Symp. on User Interface Software and Technology. 51 - 60.
- Putz-Anderson, V. (1988). Cumulative trauma disorders: A manual for musculoskeletal diseases of the upper limbs. Bristol, PA, Taylor & Francis.
- Rekimoto, J. (1996). Tilting Operations for Small Screen Interfaces. UIST'96. 167-168.
- Rekimoto, J. (1997). Pick-and-Drop: A Direct Manipulation Technique for Multiple Computer Environments. Proc. ACM UIST'97 Symp. on User Interface Software & Technology. 31-39.
- Rekimoto, J. (1998). A Multiple Device Approach for Supporting Whiteboard-based Interactions. CHI'98. 344-351.
- Rime, B. & L. Schiaratura (1991). Gesture and speech. Fundamentals of Nonverbal Behaviour. New York, Press Syndacate of the University of Cambridge: 239--281.

- Robertson, G., M. Czerwinski, K. Larson, D. Robbins, D. Thiel & M. van Dantzich (1998). Data Mountain: Using Spatial Memory for Document Management. UIST'98.
- Robertson, G., M. van Dantzich, D. Robbins, M. Czerwinski, K. Hinckley, K. Ridsen, V. Gorokhovskiy & D. Thiel (1999). The Task Gallery: A 3D Window Manager. to appear in ACM CHI 2000.
- Robertson, G. G., S. K. Card & J. D. Mackinlay (1989). The Cognitive Coprocessor Architecture for Interactive User Interfaces. Proc. UIST'89 Symposium on User Interface Software and Technology. 10-18.
- Robertson, P. K. (1988). "Perceptual color spaces. Visualizing color gamuts: A user interface for the effective use of perceptual color spaces in data display." IEEE Comput. Graphics Appl. **8**(5): 50-64.
- Rutledge, J. & T. Selker (1990). Force-to-Motion Functions for Pointing. Proc. of Interact '90: The IFIP Conf. on Human-Computer Interaction. 701-706.
- Saffo, P. (1997). Sensors: The Next Wave of Infotech Innovation. Institute for the Future: 1997 Ten-Year Forecast. 115-122.
- Sawhney, N. & C. M. Schmandt (2000). "Nomadic Radio: Speech and Audio Interaction for Contextual Messaging in Nomadic Environments." ACM Transactions on Computer-Human Interaction **7**(3): 353-383.
- Schilit, B. N., N. I. Adams & R. Want (1994). Context-Aware Computing Applications. Proc. IEEE Workshop on Mobile Computing Systems and Applications, Santa Cruz, CA, IEEE Computer Society. 85-90.
- Schmandt, C. M. (1983). "Spatial Input/Display Correspondence in a Stereoscopic Computer Graphic Work Station." Computer Graphics (Proc. ACM SIGGRAPH '83) **17**(3): 253-262.
- Schmandt, C. M., N. Marmasse, S. Marti, N. Sawhney & S. Wheeler (2000). "Everywhere messaging." IBM Systems Journal **39**(3&4).
- Schmidt, A., Aidoo, K., Takaluoma, A., Tuomela, U., Van Laerhove, K., Van de Velde, W. (1999). Advanced Interaction in Context. Handheld and Ubiquitous Computing (HUC'99), Springer-Verlag. 89-101.
- Sears, A. (1993). "Investigating touchscreen typing: the effect of keyboard size on typing speed." Behaviour & Information Technology **12**(1): 17-22.
- Sears, A., C. Plaisant & B. Shneiderman (1992). A New Era for High Precision Touchscreens. Advances in Human-Computer Interaction. Hartson and Hix, Ablex Publishers. **3**: 1-33.
- Sears, A. & B. Shneiderman (1991). "High Precision Touchscreens: Design Strategies and Comparisons with a Mouse." International Journal of Man-Machine Studies **34**(4): 593-613.
- Sellen, A., G. Kurtenbach & W. Buxton (1992). "The Prevention of Mode Errors through Sensory Feedback." Human Computer Interaction **7**(2): 141-164.

- Serra, L., N. Hern, C. Beng Choon & T. Poston (1997). Interactive Vessel Tracing in Volume Data. ACM/SIGGRAPH Symposium on Interactive 3D Graphics, Providence, R. I., ACM, New York. 131-137.
- Sheridan, T. B. (1992). Telerobotics, Automation, and Human Supervisory Control. Cambridge, MA, MIT Press.
- Sigoma, K. B. (1993). A survey of perceptual feedback issues in dexterous telemanipulation: part I. finger force feedback. Proc. IEEE Virtual Reality Annu. Int. Symp. 263-270.
- Silverstein, D. (1977). Human factors for color display systems. Color and the Computer: Concepts, Methods and Research, Academic Press: 27-61.
- Simpson, C. A. & K. Marchionda-Frost (1984). "Synthesized speech rate and pitch effects on intelligibility of warning messages for pilots." Human Factors **26**: 509-517.
- Smailagic, A. & D. Siewiorek (1996). "Modalities of Interaction with CMU Wearable Computers." IEEE Personal Communications(Feb.): 14-25.
- Smith, R. B. & A. Taivalsaari (1999). "Generalized and Stationary Scrolling." Proc. UIST'99: CHI Letters **1**(1): 1-9.
- Snibbe, S. & K. MacLean (2001). "Haptic Techniques for Media Control." CHI Letters (Proc. UIST 2001) **3**(2): 199-208.
- Spiker, A., S. Rogers & J. Cicinelli (1985). Selecting color codes for a computer-generated topographic map based on perception experiments and functional requirements. Proc. 3rd Symp. Aviation Psychology, Ohio State University, Dept. of Aviation, Columbus, OH. 151-158.
- Stifelman, L. (1996). Augmenting Real-World Objects: A Paper-Based Audio Notebook. CHI'96 Conference Companion. 199-200.
- Stokes, A., C. Wickens & K. Kite (1990). Display Technology: Human Factors Concepts. Warrendale, PA, SAE.
- Stone, M. C., W. B. Cowan & J. C. Beatty (1988). "Color gamut mapping and the printing of digital color images." ACM Trans. Graphics **7**(4): 249-292.
- Stratton, G. (1897). "Vision without inversion of the retinal image." Psychological Review **4**: 360-361.
- Streitz, N. A., J. Geißler, T. Holmer, S. Konomi, C. Müller-Tomfelde, W. Reischl, P. Rexroth, R. P. Seitz & Steinmetz (1999). i-LAND: An interactive Landscape for Creativity and Innovation. ACM CHI'99 Conf. on Human Factors in Computing Systems, Pittsburgh, PA. 120-127.
- Sugiura, A. & Y. Koseki (1998). A User Interface Using Fingerprint Recognition - Holding Commands and Data Objects on Fingers. UIST'98 Symp. on User Interface Software & Technology. 71-79.
- Sutherland, I. E. (1968). A Head-mounted Three Dimensional Display. Proc. the Fall Joint Computer Conference. 757-764.
- Swaminathan, K. & S. Sato (1997). "Interaction Design for Large Displays." interactions(Jan-Feb).

- Tan, D. S., J. K. Stefanucci, D. R. Proffitt & R. Pausch (2001). The Infocockpit: Providing Location and Place to Aid Human Memory. Workshop on Perceptive User Interfaces, Orlando, FL.
- Tan, D. S., J. K. Stefanucci, D. R. Proffitt & R. Pausch (2002). Kinesthesia Aids Human Memory. CHI 2002 Extended Abstracts, Minneapolis, MN.
- Tandler, P., T. Prante, C. Müller-Tomfelde, N. A. Streitz & R. Steinmetz (2001). Connectables: dynamic coupling of displays for the flexible creation of shared workspaces. UIST 2001. 11-20.
- Tani, M., K. Yamaashi, K. Tanikoshi, M. Futakawa & S. Tanifuji (1992). Object-Oriented Video: Interaction with Real-World Objects through Live Video. Proceedings of ACM CHI'92 Conference on Human Factors in Computing Systems. 593-598, 711-712.
- Trevor, J., D. M. Hilbert, B. N. Schilit & T. K. Koh (2001). From Desktop to Phonetop: A UI for Web Interaction on Very Small Devices. Proc. UIST '01 Symp. on User Interface Software and Technology. 121-130.
- Triesman, A. (1985). "Preattentive processing in vision." Comput. Vision, Graphics and Image Process. **31**: 156-177.
- Tufte, E. R. (1983). The Visual Display of Quantitative Information. P.O. Box 430, Cheshire, CT 06410, Graphics Press.
- Tufte, E. R. (1990). Envisioning Information. P.O. Box 430, Cheshire, CT 06410, Graphics Press.
- Tufte, E. R. (1997). Visual Explanations: Images and Quantities, Evidence and Narrative. Cheshire, CT, Graphics Press.
- Wang, J., S. Zhai & H. Su (2001). Chinese input with keyboard and eye-tracking: an anatomical study. Proceedings of the SIGCHI conference on Human factors in computing systems.
- Want, R. & G. Borriello (2000). "Survey on Information Appliances." IEEE Personal Communications(May/June): 24-31.
- Want, R., K. P. Fishkin, A. Gujar & B. L. Harrison (1999). Bridging physical and virtual worlds with electronic tags. Proc. ACM CHI'99 Conf. on Human Factors in Computing Systems. 370-377.
- Ware, C. (1988). "Color sequences for univariate maps: theory, experiments, and principles." IEEE Comput. Graphics Appl. **8**(5): 41-49.
- Ware, C. (2000). Information visualization: design for perception. San Francisco, Morgan Kaufmann.
- Ware, C., K. Arthur & K. S. Booth (1993). Fish Tank Virtual Reality. Proceedings of ACM INTERCHI'93 Conference on Human Factors in Computing Systems. 37-41.
- Ware, C. & D. R. Jessome (1988). "Using the Bat: A Six-Dimensional Mouse for Object Placement." IEEE Computer Graphics and Applications(November 1988): 65-70.

- Ware, C. & J. Rose (1999). "Rotating virtual objects with real handles." ACM Transactions on CHI **6**(2): 162-180.
- Weiser, M. (1991). "The Computer for the 21st Century." Scientific American(September): 94-104.
- Wellner, P. (1993). "Interacting with Paper on the DigitalDesk." Communications of the ACM **36**(7): 87-97.
- Wenzel, E. M. (1992). "Localization in virtual acoustic displays." Presence **1**(1): 80-107.
- Westheimer, G. (1979). "Cooperative neural processes involved in stereoscopic acuity." Exp. Brain Res. **36**: 585-597.
- Wickens, C. (1992). Engineering Psychology and Human Performance. New York, HarperCollins.
- Wilson, A. & S. Shafer (2003). XWand: UI for Intelligent Spaces. CHI 2003. to appear.
- Wilson, F. R. (1998). The Hand: How its use shapes the brain, language, and human culture. New York, Pantheon Books.
- Wisneski, C., H. Ishii, A. Dahley, M. Gorbet, S. Brave, B. Ullmer & P. Yarin (1998). "Ambient Displays: Turning Architectural Space into an Interface between People and Digital Information." Lecture Notes in Computer Science **1370**: 22-32.
- Wyszecki, G. & W. S. Styles (1982). Color Science, 2nd ed., Wiley, New York.
- Zeleznik, R., K. Herndon & J. Hughes (1996). SKETCH: An Interface for Sketching 3D Scenes. Proceedings of SIGGRAPH 96, New Orleans, Louisiana. 163-170.
- Zhai, S. (1998). "User Performance in Relation to 3D Input Device Design." Computer Graphics **32**(8): 50-54.
- Zhai, S., S. Conversy, M. Beaudouin-Lafon & Y. Guiard (2003). Human On-line Response to Target Expansion. CHI 2003 Conf. on Human Factors in Computing Systems, Ft. Lauderdale, FL. 177-184.
- Zhai, S., M. Hunter & B. A. Smith (2000). "The Metropolis Keyboard- An Exploration of Quantitative Techniques for Virtual Keyboard Design." CHI Letters **2**(2): 119-128.
- Zhai, S. & P. Milgram (1993). Human Performance Evaluation of Isometric and Elastic Rate Controllers in a 6DoF Tracking Task. Proc. SPIE Telemanipulator Technology.
- Zhai, S., P. Milgram & W. Buxton (1996). The Influence of Muscle Groups on Performance of Multiple Degree-of-Freedom Input. Proceedings of CHI'96: ACM Conference on Human Factors in Computing Systems, Vancouver, British Columbia, Canada, ACM, New York. 308-315.
- Zhai, S., C. Morimoto & S. Ihde (1999). Manual and Gaze Input Cascaded (MAGIC) Pointing. Proc. ACM CHI'99 Conf. on Human Factors in Computing Systems. 246-253.

Zhai, S., B. A. Smith & T. Selker (1997). Improving Browsing Performance: A study of four input devices for scrolling and pointing tasks. Proc. INTERACT97: The Sixth IFIP Conf. on Human-Computer Interaction. 286-292.