Computers in HCI

Robert J.K. Jacob

Department of Electrical Engineering and Computer Science
Tufts University
Medford, Mass. USA

The previous section of the handbook described the *human* in human-computer interaction; here we consider the *computer*. The problem of human-computer interaction can be viewed as two powerful information processors, the human and the computer, attempting to communicate with each other via a narrow-bandwidth, highly constrained interface[7]. However, while the computer is evolving rapidly, the human is not. If we can predict anything about human-computer interaction 100—or 1000—years from now it is that the humans will be quite similar to ourselves in size, shape, perceptual abilities, and mental capacities, while the computers might well be unrecognizable to us. As designers and builders of user interfaces, we must thus invent systems that fit the unchanging characteristics of humans, rather than to try to bend the humans to fit our changing designs. Human-computer communication was once conducted via switches, knobs, jumper wires on patch boards, and status lights. These were convenient for the computer but awkward and difficult for humans; they required considerable training and adaptation; and computers were used by only a few, highly trained engineers. Now, as computer use becomes more widespread and computer power increases, we can spend more and more of that power on better communication with the user.

Today, the frontiers of research in human-computer communication seek new ways to increase the bandwidth and naturalness of this communication channel. The challenge is to design new interaction media that better fit and exploit the characteristics of humans that are relevant to communication. We can view the fundamental task of human-computer interaction as moving information between the brain of the user and the computer. Progress in this area attempts to increase the useful bandwidth across that interface by seeking faster, more natural, and more convenient communication means. On the user's side of the communication channel, the interface is constrained by the nature of human communication organs and abilities; on the computer side, it is constrained only by the devices and methods that we can invent.

The chapters in this section describe the state of the art and the frontiers of research in this area. They focus on the actual input and output devices that we use to connect human and computer but take us well beyond the familiar mouse, keyboard, and video display. All of our designs and plans for new user experiences and human-computer interfaces, from high-level conceptual and organizational issues to low-level mechanical ones are eventually funneled through the actual input and output computer media covered here. They are the ultimate means by which the human and computer communicate information.

Hinckley surveys the world of input devices and lays out the field as a cornerstone of this section. We typically think of a mouse, keyboard, and maybe a virtual reality glove or tracker, but Hinckley places these and other devices into a broader perspective, and provides a taxonomy for understanding the range of existing and possible input devices.

Karat, Vergo, and Nahamoo cover speech both as an input and output medium. They show how the problem of speech recognition goes beyond simply processing the input and requires understanding of natural language, dialogue, and the context of the application. They also discuss multi-mode dialogues, which combine speech with gestures and other actions to make a dialogue with a computer that is more like one with other humans.

The principal means for humans to receive information is through vision. Unlike the earliest, text-based computer interfaces, our current interfaces reflect this and rely heavily on visual communication. Luczak, Roetting, and Oehme go beyond the familiar CRT and LCD displays and provide a systematic overview of the full range of visual displays. They cover the principles of display systems and survey the current technology in this field.

Another form of computer output is through the sense of touch—haptic output. Computers can generate visual images of highly complex and interactive objects and worlds, but we can typically only see them, not touch them. Iwata introduces the field of haptic output, in which computers synthesize physical sensations, much as they already do visual ones. He introduces the sensory mechanisms involved, both proprioceptive for the motion of the limbs and cutaneous for the direct sense of touch. He then describes a range of technologies for producing these sensations, some rather more exotic and less familiar than our visual display devices.

The final communication mode in our section is auditory output. We have already covered speech, but here we are interested in the use of other sounds in a user interface. Brewster surveys the field of non-speech auditory output, ranging from realistic or representational sounds of physical objects to abstract, often more musical sounds. He shows how visual and auditory output have opposite characteristics in several respects so that combining the two can make a particularly effective interface.

Beyond these chapters, we might look toward the future of computer input and output media. It is easy to predict that computers both smaller and larger than today's desktop workstations will be widespread. This will be a force driving the design and adoption of future interface mechanisms. Small computers such as laptops, palmtops, personal digital assistants, and wearable computers are often intended to blend more closely into the user's other daily activities. They will certainly require smaller input/output mechanisms, and may also require more unobtrusive devices, if they are to be used in settings where the user is simultaneously engaged in other tasks, such as talking to people or repairing a piece of machinery. At the same time, computers will be getting larger. As display technology improves, as more of the tasks one does become computer-based, and as people working in groups use computers for collaborative work, an office-sized computer can be envisioned, with a display as large as a desk or wall (and resolution approaching that of a paper desk). Such a computer leaves considerable freedom for user interaction. If it is a large, fixed installation, then it could accommodate a special-purpose console or ''cockpit'' for high-performance interaction. Or the large display might be fixed, but the user or users move about the room, interacting with each other and with other objects in the room, using small, mobile interaction devices.

A broader trend is toward increased naturalness in user interfaces. Such interfaces seek to make the user's actions as close as possible to the user's thoughts that motivated those actions, that is, to reduce the ''Gulf of Execution'' described by Hutchins, Hollan, and Norman[1], the gap between the user's intentions and the actions necessary to communicate them into the computer. Such natural interfaces build on the equipment and skills humans have already acquired and exploit them for human-computer interaction. A prime reason for the success of direct manipulation interfaces[6] is that they draw on analogies to existing human skills (pointing, grabbing, moving objects in space), rather than trained behaviors. Virtual reality interfaces, too, gain their strength by exploiting the user's pre-existing abilities and expectations.

Navigating through a conventional computer system requires a set of learned, unnatural commands, such as keywords to be typed in, or function keys to be pressed. Navigating through a virtual reality system exploits the user's existing, natural ''navigational commands,'' such as positioning his or her head and eyes, turning his or her body, or walking toward something of interest. Tangible user interfaces and ubiquitous computing similarly leverage real-world manipulation of real physical objects to provide a more natural interface[2]. The result in each case is to make interacting with the computer is more ''natural,'' more like interacting with the rest of the world and, incidentally, to make the computer more invisible.

Another trend in human-computer interaction is toward lightweight, non-command, or passive interactions, which seek inputs from context and from physiological or behavioral measures. These provide information from a user without explicit action on his or her part. For example, behavioral measurements can be made from changes the user's typing speed, general response speed, manner of moving the cursor, frequency of low-level errors, or other patterns of use. Passive measurements of the user's state may also be made with, for example, three-dimensional position tracking, eye tracking, computer vision, and physiological sensors. They can then subtly modify the computer's dialogue with its user to match the user's state[5]. Blood pressure, heart rate variability, respiration rate, eye movement and pupil diameter, galvanic skin response (the electrical resistance of the skin), and electro-encephalograms can be used in this way. They are relatively easy and comfortable to measure, but interpreting them accurately and instantaneously within a user-computer dialogue is much more difficult. We might thus look into the future of interaction devices by examining a progression that begins with experimental devices used to measure some physical attribute of a person in laboratory studies. As such devices become more robust, they may be used as practical medical instruments outside the laboratory. As they become convenient, non-invasive, and inexpensive, they may find use as future computer input devices. The eye tracker is such an example[3]; other physiological monitoring devices may also follow this progression.

It is also interesting to note that, with the current state of the art, computer input and output are somewhat asymmetric. The bandwidth from computer to user is typically far greater than that from user to computer. Graphics, animations, audio, and other media can output large amounts of information rapidly, but there are hardly any means of inputting comparably large amounts of information from the user. Humans can receive visual images with very high bandwidth, but we are not very good at generating them. We can generate higher bandwidth with speech and gesture, but computers are not yet adept at interpreting these. User-computer dialogues are thus typically one-sided. Using lightweight, passive measurements as input media can help redress this imbalance by obtaining more data from the user conveniently and rapidly.

Perhaps the final frontier in user input and output devices will someday be to measure and stimulate neurons directly, rather than relying on the body's transducers[4]. Our goal is to move information between the brain of the user and the computer; current methods all require a physical object as an intermediary. Reducing or eliminating the physical intermediary could reduce the Gulf of Execution and improve the effectiveness of the human-computer communication, to lead to a view of the computer as a sort of mental prosthesis, where the explicit human-computer communication actions vanish and the communication is direct, from brain to computer.

Acknowledgments

References

1. E.L. Hutchins, J.D. Hollan, and D.A. Norman, ''Direct Manipulation Interfaces,'' in *User Centered System Design: New Perspectives on Human-computer Interaction*, ed. by D.A. Norman and S.W. Draper, pp. 87-124, Lawrence Erlbaum, Hillsdale, N.J., 1986.

2. H. Ishii and B. Ullmer, ''Tangible Bits: Towards Seamless Interfaces between People, Bits, and Atoms,'' *Proc. ACM CHI'97 Human Factors in Computing Systems Conference*, pp. 234-241, Addison-Wesley/ACM Press, 1997.

3. R.J.K. Jacob, ''Eye Movement-Based Human-Computer Interaction Techniques: Toward Non-Command Interfaces,'' in *Advances in Human-Computer Interaction, Vol. 4*, ed. by H.R. Hartson and D. Hix, pp. 151-190, Ablex Publishing Co., Norwood, N.J., 1993. http://www.eecs.tufts.edu/˜jacob/papers/hartson.txt [ASCII]; http://www.eecs.tufts.edu/˜jacob/papers/hartson.ps [Postscript].

4. M. Moore, P. Kennedy, E. Mynatt, and J. Mankoff, ''Nudge and Shove: Frequency Thresholding for Navigation in Direct Brain-computer Interfaces,'' *Proc. ACM CHI 2001 Human Factors in Computing Systems Conference Extended Abstracts*, pp. 361-362, ACM Press, 2001.

5. R.W. Picard, *Affective Computing,* MIT Press, Cambridge, Mass., 1997.

6. B. Shneiderman, ''Direct Manipulation: A Step Beyond Programming Languages,'' *IEEE Computer*, vol. 16, no. 8, pp. 57-69, 1983.

7. E.R. Tufte, ''Visual Design of the User Interface,'' IBM Corporation, Armonk, N.Y., 1989.