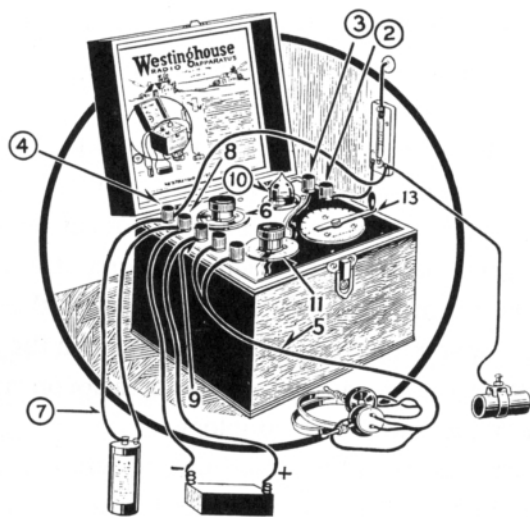


OPERATING INSTRUCTIONS FOR AERIOLA SR.



Text numbers correspond with above diagram.

Numbers Corresponding to Diagram

- No. 1. First, refer to accompanying sketch, then erect antenna and place protective device in position as described on page 56.
- No. 2. Connect a wire leading from terminal marked R on protective device to binding post indicated by arrow for stations below 350 meters.
- No. 3. For stations between 350 and 500 meters, connect the above wire to this post.
- No. 4. Connect this post with terminal G of protective device.
- No. 5. Connect telephone receivers to these two posts.
- No. 6. Turn rheostat as far as it will go toward tail of arrow.

- No. 7. Connect to positive (center) terminal of the single 1.5 volt dry cell.
- No. 8. Connect to negative (outside) terminal of the single 1.5 volt dry cell and negative terminal (-) of 22.5 volt plate battery.
- No. 9. Connect to positive terminal marked (+) of 22.5 volt plate battery.
- No. 10. Insert Aeriotron Vacuum tube in receptacle provided. Note that the four holes in base which receive prongs of tube are not all alike, one being larger than the rest, thus permitting insertion of tube in but one way. Be sure prongs register with holes and then press in firmly.
- No. 11. Place "Tickler" pointer at zero point of scale.
- No. 12. Turn rheostat (6) toward point of arrow until vacuum tube shows dull red. Do not try to burn too brightly as this materially reduces the life of the filament.
- No. 13. Rotate tuning handle slowly over the scale, meanwhile listening until sound is heard in the telephone receivers. Adjust to best position, then increase "Tickler" (11) until maximum strength of signal is obtained. If tickler is turned too far toward maximum position, signals will lose their natural tone and reception of telephone signals may become difficult.

Note: This terminal is also connected to terminal G of the protective device.

Complete Aeriola Sr. Broadcasting Receiver, Model RF, 190-500 Meters, with One Aeriotron WD-11-D Vacuum Tube, One Filament Dry Cell, One Plate Dry Battery, Head Telephone Receivers, Antenna Equipment and Full Instructions \$75.90

Aeriola Sr. Broadcasting Receiver, Model RF, As Above, Less Batteries and Antenna Equipment \$65.00

Dimensions: 7 in. x 8 1/2 in. x 7 1/4 in.

Weights: Net, 6 lbs.; Shipping, 12 lbs.; with Antenna Equipment and Batteries, 25 lbs.

NOTE: For Prices of other Complete Receiver Combinations, see page 35.

8

Why Is Everything So Difficult to Use?

"Everything that can be invented has been invented."

— Charles H. Duell, Commissioner, U.S. Office of Patents, 1899

We are fond of believing that today's world of technology is special, that the rapid pace of change, the fights among companies in the battle for supremacy, and, especially, the story of the personal computer and telecommunications revolution mark this as a special time in history. Well, sort of. If this is a special time in history, it is because every time is special. Every era is unique, having its own particular quirks and foibles. Today's technological revolution is sparked by the convergence of several industries, especially information technology, communications, and entertainment, but even here, what is happening today follows a century-old historical pattern.

In many ways, the changes during the latter part of the nineteenth century and the start of the twentieth were far more dramatic and far-reaching than what we are currently witnessing at the start of the twenty-first century. In the early 1900s, there were radical changes in lifestyles. Rapid travel was not possible before; now it was starting to be commonplace, first by means of the train and the ocean liner, then with the automobile and airplane. Talking to someone at a distance had not been possible; now, with the invention and deployment of the telephone, it was. Instant records of events, both the sights (via photography) and the sounds (via the phonograph) never were possible, now they were.

Figure 8.1

The first radios were not easy to use. From *Radio Enters the Home*. (New York: Radio Corporation of America, 1922.)

Today, the “revolution” in which we live consists mainly of improvements in what has already existed. These, of course, are important; they are changing the fabric of everyday life and the structure of society and government. But compared to the dramatic changes at the start of the twentieth century, they seem more incremental than revolutionary. Today we are amused by the statement made in 1899 by the Commissioner of the U.S. Office of Patents that “everything that can be invented has been invented,” but given the remarkable list of accomplishments and discoveries in that era, we should also feel sympathetic.

If the late 1800s and early 1900s was an era of dramatic revolution in technology, it was also the era of increased complexity of life and of a regimented work style.¹ Henry Ford’s assembly line philosophy (“Fordism”) and Frederick Taylor’s principles of scientific management (“Taylorism”) specified work practices throughout the world.² They led to the dehumanization of work, although this trend had started prior to Ford and Taylor in the textile factories of both England and New England. Feeding, clothing, and entertaining the people of the world is a huge enterprise, one in which the foot soldiers (factory and field workers) are often sacrificed. For the steadily increasing middle class and the upper class, it was a great time to be alive.

Today, we are fond of making similar claims. We live in interesting times. Information technology has reached its prime, or, at least, so we think. We sit in front of personal computers and communicate with the world. In a few seconds I can do research on the history of technology by examining locations around the world: Europe, the United States, Canada, Australia. Indeed, I am not always sure where the information resides that I am examining, and often I don’t care. National boundaries do not matter. Today, two technologies provide a powerful infrastructure: communications and computation. Put them together and it changes the face of personal interaction, business, education, and government. The financial and business sectors are breaking away from governmental boundaries and controls. The world is changing in ways difficult to forecast.

The Double-Edged Sword of Technology

The double-edged sword is a technology, and technology is like a double-edged sword. It can enhance and diminish our lives. I have long been a fan of technology—even when I was a professor and long before I worked for the computer industry. In *Things That Make Us Smart* I made the case for technology, arguing that properly built, properly deployed, technology makes people smarter than we would otherwise be. The problem lies in that word “properly.” Modern-day technology enslaves us as much as it empowers us.

The modern-day culprit is the computer, but I can make the point even better by looking at the telephone. When the telephone was first invented, in the late 1800s, everyone knew it was important, but nobody knew why. “Every city in the United States will have to have one,” the pundits said. The thought was that people would gather round the phone in the town square and listen to news and concerts. Little did they realize that it would change the fabric of business and home. In the early days, great concern was given to privacy, to the question of who would be allowed to use the telephone. One telephone company took the phone away from a hotel: why, it was allowing mere guests to use it. What would happen, the telephone company asked, if anyone could simply call up anyone else? Horrors.

Those early concerns were soon brushed away; today they seem quaint. But perhaps they were also correct. The telephone today is abrasive and intrusive. A caller has no way of knowing whether the person being called is busy or idle, in a good mood or foul. The person getting the phone call has no way of knowing if the message is important or a bother, whether it will be a short conversation or a long one.

The telephone answering machine has improved life to some extent. It allows us to decide whether we wish to answer. The answering machine itself is a double-edged technology, one with an interesting history. When it first appeared, it was mainly used by business, in particular small businesses without anyone who could answer the

phone and take messages when the intended recipient was not there. When the first answering machines appeared in homes, most callers considered them to be rude; people were annoyed to get an automated response on a machine when what they wanted was a person.

Today, at least in many countries of the world, the answering machine is taken for granted. In some places, it is considered rude not to have such a machine, for then it isn't possible to leave messages when the person is either away from the phone or otherwise not answering. Not only that, but there are times when we simply wish to convey a short message without a lengthy conversation. In this case, we sometimes prefer just to get the answering machine. How many times have you called someone and been annoyed to get a person rather than the machine, especially when the person who answered the phone was not the one for whom the message was intended? "Let me call back and just leave a message on the machine," you request of the person.

The transition of the answering machine from that of a rude, inconsiderate technology to the class of an essential, convenient one, is similar to the path followed by other devices. Most technologies go through a cycle of initial rejection, followed by experimentation and transformation as their true value is discovered. The original answering machines played aloud the caller's voice as their message was being recorded. I don't know if this was purposeful or not, but surely the original designers of these machines—who had business users in mind—could not have anticipated that this feature would be used in homes to screen calls, to decide if the caller was somebody anyone at the receiving end wished to speak with at that moment.

But there is more to come. Wireless telephones and pagers are increasingly popular, not just among business people but also in the home, especially with teenagers. Today, some information technologies are so important that they get fastened to the body: the watch, eyeglasses, and hearing aids. Soon, the telephone will join the watch, always with us, always available. But what then of privacy, of peace and quiet? Of time to ourselves? The tremendous convenience of continual access carries with it an invasion of privacy and personal space.

Will the telephone really join the watch? I think it will supersede the watch and join the pacemaker. The watch has become jewelry, worn on the wrist as much for fashion's sake as for telling the time. But a phone? Why not permanently implant it in the head? After all, wiggle the skin just below the ear, at the jawbone. Feel that extra slack? Just enough to make a tiny cut, implant a tiny telephone, put the microphone and earpiece on the bone near the ear, and there you have it: a telephone always available, no matter where you are, no matter what you are doing. Every so often it will be wonderful. Every so often it will be a nightmare come true.

Complexity and Difficulty

There are two faces to the complexity of a device, one internal to the machinery, the other external to the world and to the user. The complexity that concerns me is the second kind: the external complexity that determines how easy or difficult the device is to use. To distinguish between these two meanings of complexity, let me use *complexity* to refer to the internal mechanisms of the device and *difficulty* to refer to the factors that affect ease of use.

Along with advances in technology come the dreaded curses of technology: difficulty in operation and diminished control over our lives. Technology can be confusing, maddening, infuriating. Technology is its own master. It has its own requirements and own rules of operation. The complexity, difficulty, and demands of technology have long been a source of complaint. That people must conform to technology seems to be the basic premise, and the notion that it is technology that should conform is a modern invention. In many ways, the conformance of technology wasn't even possible until recently, for up until recent times, it was remarkable enough that the complex mechanical and electrical devices would work at all, let alone be designed to be accessible and usable by the average person. This is the reason for the paradox that today's technology is largely built from a machine-centered point of view even though it is designed and built by humans.

Early technology was not only difficult, it was dangerous. Missing fingers, even thumbs, are still common among farmers, construction workers, and miners. At the start of the industrial revolution,³ hunger and poverty were prevalent. Disease was everywhere, and plagues could kill entire communities. The working conditions within many industries were hazardous, and there was slight regard for safety—that was the responsibility of the worker. Clothes could be caught in machinery, resulting in death or serious injury. It is only in recent times that we have started to care for the welfare of the workers and put the blame on the design or the work procedures rather than on the worker.

It is still common to hear managers say that if someone gets injured on the job, that individual must have been doing something wrong. This is the “blame and train” philosophy. This philosophy makes the managers’ lives easier, and it is one with which even the injured workers will often agree.

“That was stupid of me,” they say, shaking their heads. “How did I stick my finger into that fan blade?” the worker says. “I knew it was there, I knew it was dangerous.”

True, but why was the fan blade designed so that a finger could be stuck into it? Why was it located where the finger could reach? Just because people are willing to blame themselves for their injuries does not mean they are correct to do so.

The new technologies are designed to be used by people, ordinary people, people who grow fatigued, whose attention wanders, whose mind is preoccupied. It does no good to legislate against such properties of human nature. It does no good to complain that if only workers would keep their minds focused on the task, they would not be getting injured. Everyone’s mind wanders, everyone daydreams, gets fatigued, workers and management alike. Proper design takes this into account.

A great deal of effort is aimed at preventing analogous problems in machines. Engineers and designers take account of metal fatigue and random electrical noise. They need to do the same for human functioning. It’s not easy. The “blame and train” philosophy seems deeply ingrained within our consciousness, in part a relic of the Fordism and

Taylorism of the early 1900s. This is a philosophy that avoids the true source of the problem: badly designed technology, badly designed procedures.

The dilemma has been with us a long time, and I suspect it will be with us yet for a long time. This is yet another legacy of the era of scientific management, where the human is treated as a machine and then found to be deficient. Someday the values might change, letting it be acknowledged that we are complementary to machines. When we evaluate human skills and abilities according to human values, it is people who are found to be superior, machines deficient.

The Origins of Technological Difficulty

Others have speculated on why today’s technology seems so much more difficult and complex than that of earlier years. One author thought it was due to covering up the working parts. He suggested that:

If the discovery of electricity changed the face of the earth, it also altered the look of every tool, appliance, or piece of machinery that it touched. You cannot see the workings of any electrical device. All that mysterious stuff is placed within a housing of some sort—out of reach, out of sight, and beyond control. This 1905 chopper shows the way things used to be. The works, blades, gears are all there for all the world to behold. People had to know how things worked, so when things stopped working, they could be fixed.⁴

There is much to be said for this idea, but it isn’t entirely true. Many mechanical devices were difficult to use even though (or perhaps because) all of the moving parts were visible. Electricity is not the culprit. Difficulty seems to go hand-in-hand with technology, from early mechanical devices through today’s information-based ones.

How long has technology been so difficult? Perhaps forever. The joint evils of poor manuals and overcomplexity resulting from “creeping featurism” attacked the common farm tool, the plough, almost 500 years ago.⁵ Where is the difficulty in something as simple as a plough? The problem was that those sixteenth-century technologists kept adding features and adjustments. There was a coulter to cut roots in front of the share and a mouldboard to deflect the sod as it was cut. Sometimes there

was a wheel, depending upon the type of soil. The curvature of the mouldboard was adjustable (and it was also twisted). The blade could be wood or iron covered. The plough then had to be adjusted for the kind of soil, the amount of water (muddy versus dry soil), the condition of the soil, the amount of existing vegetation, and the kind of plant being prepared for.

The real boom in difficulty came with the industrial revolution as continual improvements in farming, mining, manufacturing, and transportation were introduced. The late 1800s and early 1900s saw the emergence of the electrical industry, starting with communications (the telegraph, stock ticker, and telephone), along with devices for light, heat, and power (light bulbs, heating devices, and electric motors). These moved into the home and office. Today the phonograph, the telephone, and typewriter are considered simple; not so when they were first introduced.

For example, the earliest phonographs were completely mechanical, driven by a hand crank and by the acoustic energy of the source. The spring-driven and the electric motor came later, but they didn't necessarily make the device simpler. The phonograph took weeks to master and, in the end, led to the failure of the first generation of products. Early users were asked to persevere:

"If the first trial . . . is not pleasing, try it again and persevere at it. . . . stick to it. Give the Phonograph a thorough trial of two weeks."⁶

"Why should we spend time and money to learn telegraphy, shorthand and typewriting and then have an idea that no time is required to learn the Phonograph; the most delicate of all and one of the most useful.

"The use of the Phonograph must be learned, the same as anything else. . . . [One needs] a few days to learn everything about it and only a few weeks practice to acquire all the dexterity in its use."⁷

To be fair, the difficulty was to enable the phonograph both to record and play back sound, but even playing back was a trial. The point is that difficulty in the use of our technology has been with us for a long time. Many of those early technologies were difficult and, worse, dangerous.

The steam engine kept exploding, with numerous deaths and government investigations before it was tamed. The automobile was viewed with alarm. Drivers had either to be expert mechanics or be accompanied by one. Vehicles were difficult to control, and governments passed restrictive legislation. The telephone was a special technology that took time to master and, for that matter, to figure out what it was good for.

Robert Pool argues that difficulty follows almost inevitably from the need to improve.⁸ The earliest steam engines, those of Thomas Newcomen, were extremely simple, but inefficient. James Watt improved the efficiency considerably, but at the cost of added complexity. Watt's version required extra valves, parts, and critical timing. The piston had to fit more snugly into the cylinder, so the manufacturing process had to be held to much closer tolerances. The engine ran at a higher temperature—increasing the temperature of the piston was one of the means of increasing efficiency—so new lubricants and seals had to be developed that could withstand these changes. The more efficient machines produced greater power, not just in the downward direction, which is all Newcomen's machine could do, but in both the downward and upward cycles. This increased the stress upon the parts. Further efficiency required raising the pressure of the steam, which increased the danger of explosion, leading to even more complexity in the attempt to add safeguards.

Most technology goes through cycles of development and change in both internal and external complexity. Often the very first device is simple, but crude. As the device undergoes the early stages of development, its power and efficiency improve, but so does its complexity. As the technology matures, however, simpler, more effective ways of doing things are developed, and the device becomes easier to use, although usually by becoming more complex inside.

Early radios used a simple "cat's whiskers" and earphones: a metal contact (the "cat's whisker") upon a germanium crystal. With the introduction of vacuum tubes, the sets became much more powerful but required multiple controls to modulate the amplification level and tuning of each stage. Some required adjustment of filter bandwidth,

depending on what kind of signal was being received and how much background noise could be tolerated by the listener. As the illustration at the opening to this chapter shows, the instructions for operation of these early radios could be rather intimidating. The user of early radios was expected to know “the five fundamentals of radio reception”:⁹ intercepting, tuning, detecting, amplifying, and reproducing. Today’s radios are trivial to use, sometimes having only an on-off switch, a volume control, and a set of buttons to choose the station. The external simplicity comes with a greatly increased internal complexity within the electronic circuits themselves, a complexity that would astound the electrical engineers who worked on the early vacuum-tube sets. (See figure 2.1, page 22.)

A similar cycle is followed by most technology. The modern automobile is simpler for the driver than ever before, but the auto itself has greatly increased in complexity. It has thousands of parts from several technologies—mechanical, hydraulic, and electronic—so many that probably no single individual understands them all. The modern airplane is simpler and safer for the pilots and passengers than those of only a few decades ago, but it contains millions of parts that take hundreds of designers to construct. Even small devices have been modified in the years since their initial development for increased efficiency, improved performance, increased safety, smaller size, less power consumption, and less environmental impact. The result is enhanced performance and ease of use, but increased internal complexity.

I once argued that modern technology is difficult to use because its operations are invisible.¹⁰ Information is abstract, conveyed by electrons and voltage levels. The workings of modern devices are hidden inside the sealed parts. We are left to the mercy of whatever conceptual model the designer decides to impose upon us, and the result is often confusion and misunderstanding. In the old days, I argued, the user had a hope of figuring out how a thing worked by manipulating it and watching what moved. Well, I was wrong. The difficulties with the plough, with the mechanical phonograph, and with other older technologies show that this explanation is too simplistic. Some of those

early mechanical devices were extremely difficult, with such a variety of mechanisms that they were incomprehensible. The first digital computer was entirely mechanical, as were early analog computers, and they were dauntingly complex and difficult.

It is the designer’s responsibility to overcome the difficulties, to provide coherence and understanding for the user. If the designer makes the operations vague, ambiguous, and hidden from sight, provides controls with no obvious meaning, and provides little or no feedback, then the result is guaranteed to be confusing. When the user lacks a clear conceptual understanding of the device, the result is difficulty of use.

The same principles apply whether the device is mechanical or electronic. The only difference is that many mechanical devices offer the possibility of being self-explaining, of having the purpose and function visible, of offering continual and immediate feedback for each action. Electronic information devices, on the other hand, by their very nature control invisible, symbolic events, and so the possibility of a clear conceptual understanding is entirely up to the competence of the designer. When it comes to the ability of engineers to design things so that the average user can understand them, the level of competency is low indeed. This is not meant to be an indictment of engineers. After all, they are trained in engineering, not in human-centered design. To change things requires a design philosophy that focuses upon ease of use, upon providing an appropriate conceptual model, making it visible, and making all actions and displays consistent with that model.

What Makes Things Easy to Use?

When is something easy to use? My studies convince me that even the most difficult of things becomes easy when users feel they are in control, that they know what to do, when to do it, and what to expect from the device whenever they perform an operation: in other words, when they have acquired understanding. What makes something understandable? Technical knowledge is not required, just functional. What is critical is to have a good conceptual understanding of the device. Few of us

understand the technologies of the automobile, radio, or television, yet we feel comfortable with these devices because each control has a known function, we can tell when the device is working properly, and we know what to do when there are problems. We feel uncomfortable when we are out of control, when we do not know how to respond, or when our actions do not lead to the results we expect.

A feeling of control, a good conceptual model, and knowledge of what is happening are all critical to ease of use. The controls must be recognizable, it must be easy to remember their function and operation, and they must provide immediate and continual feedback about the state of the system.

When is something difficult? When the controls and actions seem arbitrary, when the system can get itself into peculiar states, peculiar in the sense that the person using it does not know what it is doing, how it got there, or how to recover. When there is a lack of understanding.

Understanding comes about when the system presents a clear conceptual model of itself, making the possible actions apparent. I called this state “knowledge in the world”¹¹—when the world itself helps tell you what to do, so no instructions, no courses, no manuals are necessary. A solid wall tells you that you can’t walk in that direction. A door signals passage. A properly constructed door even specifies whether it is to be pushed or pulled, slid or lifted, by the construction of its handles, no labels or words required. If you have to add a sign that says *push* or *pull*, then this indicates that the door is not as simple as is possible; its design is faulty.

Did you ever look at an unfamiliar tool or appliance and try to figure out what it was for and how it might work? Usually the tool makes no sense. Once you are told what it is used for, however, the previously puzzling construction of weird parts suddenly all fit together wonderfully. At first, the device presents itself as a bunch of apparently unrelated affordances. This part can stick into something, this part permits turning, this part cuts. But any object has a large number of affordances, and an assemblage of parts has a huge number of alternative combinations. What is lacking is a cohesive story to fit them together for some




Figure 8.2

Without a conceptual model, without knowing what a device is to be used for, it can appear mysterious and confusing. Know that this is an apple peeler, know that an apple is placed on the prongs and, as the handle is turned, the apple rotates under the blade, and now the structure is clear. Produce a good conceptual model, and no manual is necessary. (Photograph courtesy of Corbis-Bettmann.)

purpose—a conceptual model. In the case of the kitchen appliance, the name alone is often sufficient to provide coherence.

“That’s an apple peeler? Oh, I see. So the apple must fit here, and these prongs go into the core, holding the apple and this crank turns it. This cutting edge must be the peeler, and this crank turns the apple in front of the blade. And this is the blade that moves along the outside, cutting off the skin. Oh yes, now I see. It’s obvious.”

Even if you know what a tool is for, unless you understand how it works, its use is limited. It’s possible to use things without

understanding by just following instructions. This works until either something goes wrong or there is a need to do something unusual, something not covered by the instructions. In either case, the result can be frustrating; there is no hint of what to do next. Once there is understanding, it's possible to explain the problems and predict what to do. It's possible to invent new applications and to make corrections when things go wrong. With understanding comes a feeling of control.

Given a chance, people are wonderfully good at making sense of the world. People see faces and objects in clouds. They see patterns in tea leaves. They give explanations of people's behavior, even people they don't know. We human beings are sense-makers, making sense of the way we experience the world, but only if there is something to go on, some hints and clues as to what is happening, and why.

When I encounter a new situation, how do I know what to do? I look, listen, and copy. I try to understand what is happening. I see if I can find anything that looks familiar, and if I do, then I'll perform the actions that work in the familiar situation. Is it a new restaurant? I have to decide whether I seat myself or whether I wait for someone to seat me. If the latter, I have to decide who it is that helps me, where I should stand while waiting. To answer these questions, I look around and try to find clues. If I'm supposed to wait, is there a place that looks like a waiting area? Is there a table or stand that looks like it belongs to the head waiter, or a host who will guide me to my seat or tell me how long I must wait? I look for any clues I can find.

When I encounter a new piece of technology I do the same thing. I look at it and try to see if anything looks familiar. I look to see if there are any clear indications of what to do. And when I do something that causes a reaction, I try to understand why whatever I did led to that result.

In other words, in new situations we look for familiar patterns, we look for any signs that might direct us, and we try to make sense of whatever happens. In general, people make up explanations, stories of events that help us make our way through novel and complex situations until they become understandable and comfortable.

The Conceptual Model

The use of a good conceptual model is so fundamental to good design that I would like to discuss it here in greater detail than presented in chapter 7. What does it mean to understand how something works? Do I really have to understand automobile mechanics to drive my car, or to understand solid state physics and computer programming to use my computer? Of course not. But what I do need is a good conceptual grasp of what is going on, an understanding of the different controls and alternative actions I can take and what their impact is on the device. I need a story that puts together the behavior and appearance of a device in a sensible, comprehensible pattern. Good designers present explicit conceptual models for users. If they don't, users create their own mental models, which are apt to be defective and lead them astray.

The word processor on which I am typing this chapter makes visible an excellent conceptual model for its operation. There is a ruler on the top of the page with several sliding pointers. What is particularly nice about the design is that I can experiment by moving each slider in turn and seeing what happens. As a result, I have formed a good conceptual

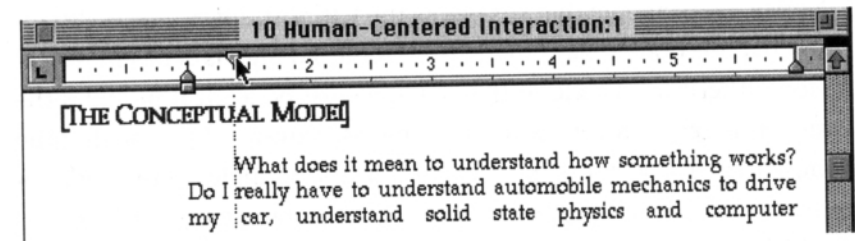


Figure 8.3
Microsoft Word's excellent conceptual model for adjusting margins. The ruler, along with the triangular and rectangular sliders are clear indicators of margin adjustments. The ruler starts and ends at the edge of the printed region of the document. Moving the mouse pointer to the downward-facing triangle and depressing the mouse button produces a vertical dotted line, helping define this triangle as the indentation for the first line of a paragraph and also aiding the user in adjusting it to the desired value. The difference between the rectangular slider and the upward-pointing triangle is not so well indicated.

model of the role the sliders play in the composition of the page. As soon as I move a slider, a dotted vertical line drops down over the text, confirming the conceptual model that the slider adjusts the location of the margins on the page, as well as letting me see just where the slider is positioned relative to the existing text. The leftmost slider sets the left margin. The next slider over, the upper slider, sets the left indentation for the first line of each paragraph (as shown in the illustration). The rightmost slider sets the right margin.

That's a useful conceptual model. I don't know how the word processor really works, deep inside its intricate programming structures. I don't know anything about all the gory programming details that couple the pointer movements to the part of the page where the action takes place. Part of my conceptual model is that the action only takes place on the line or paragraph where the cursor is located, unless I have highlighted a section of text, in which case it applies to everything that is highlighted. With this understanding I feel empowered to adjust the formatting of the page. I am in control. Notice how skillfully the graphical design conveys the model: The most clever component is the vertical dotted line that appears only when the slider is activated. It is really the vertical line that conveys the conceptual model.

Notice the rectangle that sits below the leftmost upward-facing triangle. What role does it play? Here, the graphical design is deficient. Clicking on either the triangle or the rectangle just beneath it provides the same visual result: a vertical dotted line. Moreover, if I try moving the item, whether I click on the rectangle or the triangle, both move together as a single unit. This implies that the rectangle has no function. Could it be just a design element to distinguish that first triangle from the others? A more careful examination shows that the outcome of moving the rectangle differs from that of moving the triangle. When the triangle is moved, it affects the paragraph left margin, but leaves the indentation for the first line unchanged. This lets me adjust the following lines to be left of, equal to, or right of the first line. When I move the rectangle, both the margins for the first line and the rest of the paragraph move simultaneously, as if they were locked together. This lets me widen or narrow the paragraph, leaving the amount of indentation of

the first line unchanged. Too bad the graphical design doesn't support this conceptual model.

The proper graphical design would show an explicit linkage between the rectangle and both sliders. One way of doing this would be to make the rectangle extend in length between the two controls: If I move a control, only that part of the margin is affected, and the rectangle would grow or shrink accordingly. If I moved the rectangle, it would move in position without changing its length, thereby reinforcing the conceptual model of the operation. As it is, the relationship is more difficult to discover and remember than necessary.

This point may seem like a minuscule detail, and in some respects it is. But usability often lies in the details. How many users of this word processor never discover the operation of the rectangle? Worse, how many are confused because of what they believe to be capricious results? Sometimes when they attempt to move the left margin, they change both the first line and the others, sometimes only the others. Will they be astute enough to recognize that the exact region of the slider they touched was what mattered, the triangle or the rectangle beneath it? Details matter, especially when they impact the user's conceptual model and, thereby, the user's understanding.

The basic principle is this: Start with a simple, cohesive conceptual model and use it to direct all aspects of the design. The details of implementation then flow naturally from the conceptual model.

To summarize, a conceptual model is a story. It doesn't have to discuss the actual mechanisms of the operation. But it does have to pull the actions together into a coherent whole that allows the user to feel in control, to feel there is a reason for the way things are structured, to feel that, when necessary, it's possible to invent special variations to get out of trouble and, in general, feel mastery over the device.

The job of creating a good conceptual model is in the hands of the designer. The model has to be coherent, understandable, and sufficiently cohesive that it covers the major operations of the system. It is successful if the users can tell a story, can explain to others how it all works. It is successful if the users can then use the system in ways the developers never imagined. Above all, the user should be able to discover

and learn how to use it with a minimum of effort. In the ideal case, no manual would be required. It is unsuccessful if users simply follow instructions with no understanding; they will be unable to go beyond the instructions, unable to perform novel activities, unable to get out of trouble.

Telling a coherent, consistent story that is readily understood by the wide range of users is not an easy task. It is a task best left to experts. Task-centered development, human-centered development, conceptual model-guided development: These are the secrets of success. There are more, but these are essential starting points.

Why Metaphors Should Be Avoided

A myth has grown up in the land that good interface design requires the development of a metaphor: "What metaphor shall we use to design this device?" I am often asked. "If only we would have the correct metaphor, then people would find it easy to use, right?" Wrong.

A metaphor is always wrong, by definition. After all, what do we mean when we design by means of a metaphor? We mean that we try to pick something else to guide the design, something that already exists and that we believe the users of our devices will be familiar with. The word processor is like a typewriter. The display on the page is like a piece of paper. The background of my computer screen is like a desktop, and the files and icons that are displayed upon it are like file folders, piles of paper, and physical objects on my real desk.

No, they are not. The objects on my computer screen aren't at all like the real objects with which they share a name. Metaphors are an attempt to use one thing to represent another, when the other is not the same. But if it is not the same, how can the metaphor help?

True, where the properties of the metaphor and the new thing are closely related, the metaphor helps in acquiring those properties. But when they differ, the metaphor can get in the way of learning; it either provides the wrong model or it slows up acquisition of the correct one.

Basically, those who espouse the use of metaphors are giving human-centered development a bad name, almost as bad as those who believe

in "user-friendly" systems. User-friendly systems invariably aren't. They are cloying, annoying, in-your-face systems that force the user through a sequence of steps, whether or not these are appropriate, whether or not the user is at a far more advanced level than the steps assume.

It is true that use of a metaphor is appropriate in the initial stages of learning. But while those first stages are only there temporarily, the metaphor is with us forever. After those first few steps of learning, the metaphor is guaranteed to get in the way, because by the very nature of metaphor, the thing being represented by the other isn't the same. The result is conflict. The "window" in the computer interface has a quite different meaning than the word "window" in a home. A real window does not have a "scroll bar," nor a way to change size. When the window at home is shut, it doesn't go away, out of sight. What value is added to the computer image by thinking metaphorically? For every point that the metaphor helped in understanding, it confuses at all the other points where the computer version differs from the real one.

Designers of the world: Forget the term "metaphor." Go right to the heart of the problem. Make a clean, clear, understandable conceptual model. Make sure the user can learn and understand it. Make all the actions consistent with that conceptual model. Provide feedback and interaction consistent with that model, the better to reinforce it. Forget metaphors—they will only get in the way.

Making Computers Easy to Use

Computers are general-purpose devices, designed to do everything. As a result, they can't be optimized for any individual task. Their controls are arbitrary, often limited to what can be typed on the keyboard or pointed at on the screen and clicked by the mouse. There is no possibility of using the physical construction itself to present the conceptual model of the device or of the actions, since the range of functions is so large.

With a special purpose device, everything about it can shout "Here is how I work, this is what I do." The shape of the case, the layout of the controls, even the shape and form of operation of each control indicates

something significant. The displays and feedback can be precise and to the point: explaining to the user what is going on, why the actions are needed, what they are for.

Good design is not necessarily self-explanatory: Some tasks are inherently too complex. The notion that design can render anything “intuitively obvious” is false. In fact, intuition is simply a state of subconscious knowledge that comes about after extended practice and experience. With minor exceptions, things that we call intuitive are simply skills that we have practiced for so many years that we no longer recall how difficult it was to learn them in the first place. Skills such as using a pencil, driving an automobile, speaking and understanding language, reading and writing: All these are intuitive to the skilled adult, yet all took years to learn. Difficult tasks will always have to be taught. The trick is to ensure that the technology is not part of the difficulty.

Devices for complex tasks must of themselves be complex, but they can still be easy to use if the devices are properly designed so that they fit naturally into the task. When this is done, learn the task and you know the device. One goal is that each operation fit so elegantly within the structure of the task that it need be explained only once. If users say “Of course, I see,” and never have to be told again, there is success. If users say, “Uh, I guess,” and discover that with each usage they must once again ask or consult a manual, then the design is faulty.

The goal is to let you pick up a device and use it. The goal is for the device to be built for the job and be no more difficult than it needs to be. Does this come at a cost? Of course. Among other things, it means that there must be one device per activity, which means a multitude of devices. It means that the devices are separate, so something extra must be done if you want to combine the output of several activities. All of these problematic situations can be overcome, as will be shown. The solutions require careful consideration of what things go together, of just what is meant by an activity. The solutions require a universal communication scheme so that the products of one device can be sent effortlessly to another with no more thought than if they resided on the same machine—and considering the complexity of sharing material on today’s

machines, the separate devices might even make sharing easier than today’s overloaded systems. Life is full of tradeoffs, with each choice being better for some things, worse for others.

The ultimate goal is simplicity. Make things fit the task, make the difficulty of our tools match the difficulty of the job to be done.

How do we make things less difficult? Sorry, but the causes of today’s difficulties are too fundamental for simple change. As we have seen in chapter 5, there is no magical cure that will make everything all right, letting us proceed essentially as we are now doing. What is the answer? The most promising hope is a new process for product development that focuses upon human comprehension, a human-centered development process. Read on!