CHAPTER THREE

KNOWLEDGE IN THE HEAD AND IN THE WORLD

A friend kindly let me borrow his car. Just before I was about to leave, I found a note waiting for me: "I should have mentioned that to get the key out of the ignition the car needs to be in reverse." The car needs to be in reverse! If I hadn't seen the note, I never could have figured that out. There was no visible cue in the car: the knowledge needed for this trick had to reside in the head. If the driver lacks that knowledge, the key stays in the ignition forever.

It is easy to show the faulty nature of human knowledge and memory. A common classroom exercise in the United States demonstrates that students cannot recall the pairing of letters and numbers on their telephones. One of my graduate students found that when professional typists were given caps for typewriter keys, they could not arrange them in the proper configuration.[1] American students dial telephones properly, and all those typists could type rapidly and accurately. Why the apparent discrepancy between the precision of behavior and the imprecision of knowledge? Because not all of the knowledge required for precise behavior has to be in the head. It can be distributed—partly in the head, partly in the world, and partly in the constraints of the

world. Precise behavior can emerge from imprecise knowledge for four reasons.

1. *Information is in the world.* Much of the information a person needs to do a task can reside in the world. Behavior is determined by combining the information in memory (in the head) with that in the world.

2. *Great precision is not required.* Precision, accuracy, and completeness of knowledge are seldom required. Perfect behavior will result if the knowledge describes the information or behavior sufficiently to distinguish the correct choice from all others.

3. *Natural constraints are present.* The world restricts the allowed behavior. The physical properties of objects constrain possible operations: the order in which parts can go together and the ways in which an object can be moved, picked up, or otherwise manipulated. Each object has physical features—projections, depressions, screwthreads, appendages—that limit its relationships to other objects, operations that can be performed to it, what can be attached to it, and so on.

4. *Cultural constraints are present.* In addition to natural, physical constraints, society has evolved numerous artificial conventions that govern acceptable social behavior. These cultural conventions have to be learned, but once learned they apply to a wide variety of circumstances.

Because of these natural and artificial constraints, the number of alternatives for any particular situation is reduced, as are the amount and specificity of knowledge required within human memory.

In everyday situations, behavior is determined by the combination of internal knowledge and external information and constraints. People routinely capitalize on this fact. They can minimize the amount of material they must learn or the completeness, precision, accuracy, or depth of the learning. People can deliberately organize the environment to support their behavior. Some people with brain damage can function so well that even their co-workers may not be aware of their handicap. Nonreaders have been known to fool others, even in situations where their job presumably requires reading skills. They know what is expected of them, follow the behavior of their co-workers, and set up situations so that they do not need to read or so that their co-workers do the reading for them.

What is true in these extreme cases must certainly also be true of ordinary people in ordinary situations: it is only the amount of reliance upon the external world that differs. There is a tradeoff between the amount of mental knowledge and the amount of external knowledge required in performing tasks. People are free to operate variously in allowing for this tradeoff.

Precise Behavior from Imprecise Knowledge

INFORMATION IS IN THE WORLD

Whenever information needed to do a task is readily available in the world, the need for us to learn it diminishes. For example, we lack knowledge about common coins, even though we recognize them just fine (figure 3.1). Or consider typing. Many typists have not memorized the keyboard. Usually each letter is labeled, so nontypists can hunt and peck letter by letter, relying on knowledge in the world and minimizing the time required for learning. The problem is that such typing is slow and difficult. With experience, of course, hunt-and-peck typists learn the positions of many of the letters on the keyboard, even without instruction, and typing speed increases notably, quickly surpassing handwriting speeds and, for some, reaching quite respectable rates. Peripheral vision and the feel of the keyboard provide some information about key locations. Frequently used keys become completely learned, infrequently used keys are not learned well, and the other keys are partially learned. But as long as the typist needs to watch the keyboard, the speed is limited. The knowledge is still mostly in the world, not in the head.

If a person needs to type large amounts of material regularly, further investment is worthwhile: a course, a book, or an interactive computer program. The important thing is to learn the proper placement of fingers on the keyboard, to learn to type without looking, to get knowledge about the keyboard from the world into the head. It takes several hours to learn the system and several months to become expert. But the payoff of all this effort is increased typing speed, increased accuracy, and decreased mental load and effort at the time of typing.

There is a tradeoff between speed and quality of performance and mental effort. Thus, in finding your way through a city, locating items in a store or house, or working complex machinery, the tradeoff can determine what needs to be learned. Because you know that the information is available in the environment, the information you internally code in memory need be precise enough only to sustain the quality of behavior you desire. This is one reason people can function well in their environment and still be unable to describe what they do. For example, a person can travel accurately through a city without being able to describe the route precisely.

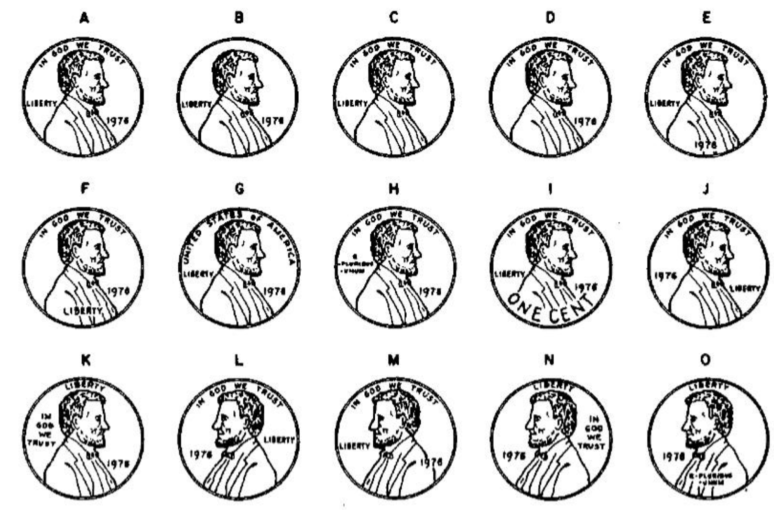

3.1 **Which Is the U. S. One Cent Coin—The Penny?** Fewer than half of the American college students who were given this set of drawings and asked to select the correct one could do so. Pretty bad performance, except that the students, of course, have no difficulty using the money: in normal life, we have to distinguish between the penny and other U.S. coins, not between several versions of one denomination. (From Nickerson & Adams, *Cognitive Psychology, 11,* © 1979. Reprinted by permission of Academic Press.)

People function through their use of two kinds of knowledge: knowledge *of* and knowledge *how.* Knowledge *of*—what psychologists call declarative knowledge—includes the knowledge of facts and rules. "Stop at red lights." "New York City lies on a parallel a bit south of Madrid, San Diego's longitude is east of Reno." "To get the key out of the ignition, the car must be in reverse." Declarative knowledge is easy to write down and to teach. Knowledge *how*—what psychologists call procedural knowledge—is the knowledge that enables a person to perform music, to stop a car smoothly with a flat tire on an icy road, to return a serve in tennis, or to move the tongue properly when saying the phrase "frightening witches." Procedural knowledge is difficult or impossible to write down and difficult to teach. It is best taught by demonstration and best learned through practice. Even the best teachers cannot usually describe what they are doing. Procedural knowledge is largely subconscious.

Knowledge from the world is usually easy to come by. Designers provide a large number of memory aids. The letters on the typewriter keyboard are one example. The lights and labels on controls act as external memory aids, reminding the user of the purpose and state of the control. Industrial equipment is replete with signal lights, indicators, and other reminders. We make extensive use of written notes. We place items in specific locations as reminders. In general, people structure the environment to provide a considerable amount of the information required for something to be remembered.

Many people organize their lives in the world, creating a pile here, a pile there, each indicating some activity to be done, some event in progress. Probably everybody uses such a strategy to some extent. Look around you at the variety of ways people structure their rooms and desks. Many styles of organization are possible, but the physical arrangement and visibility of the items frequently convey information about relative importance. Want to do your friends a nasty turn? Do them a favor—clean up their desks or rooms. Do this to some people and you can completely destroy their ability to function. [2]

GREAT PRECISION IS NOT REQUIRED

Normally, people do not need precise memory information. People can remember enough to distinguish one familiar coin from another although they may be unable to remember the faces, pictures, and words on the coins.[3] But make more precise memory necessary and you get havoc. Three countries have rediscovered this fact in recent years: the United States, when it introduced the Susan B. Anthony one-dollar coin; Great Britain, when it introduced the one-pound coin; and France, when it introduced a new ten-franc coin. The new U.S. dollar coin was confused with the existing twenty-five-cent piece (the quarter), and the British pound coin was confused with the existing five-pence piece. (The one-pound coin has the same diameter as the five-pence piece, but is considerably thicker and heavier.) Here is what happened in France:

"PARIS . . ." With a good deal of fanfare, the French government released the new 10-franc coin (worth a little more than $1.50) on Oct. 22 [1986]. The public looked at it, weighed it, and began confusing it so quickly with the half-franc coin (worth only 8 cents) that a crescendo of fury and ridicule fell on both the government and the coin.

"Five weeks later, Minister of Finance Edouard Balladur suspended circulation of the coin. Within another four weeks, he canceled it altogether.

"In retrospect, the French decision seems so foolish that it is hard to fathom how it could have been made. . . . After much study, designers came up with a silver-colored coin made of nickel and featuring a modernistic drawing by artist Joaquim Jimenez of a Gallic rooster on one side and of Marianne, the female symbol of the French republic, on the other. The coin was light, sported special ridges on its rim for easy reading by electronic vending machines and seemed tough to counterfeit.

"But the designers and bureaucrats were obviously so excited by their creation that they ignored or refused to accept the new coin's similarity to the hundreds of millions of silver-colored, nickel-based half-franc coins in circulation . . . [whose] size and weight were perilously similar."[4]

The confusions probably occurred because the users of coins formed representations in their memory systems that were sufficiently precise only to distinguish among the coins that they actually had to use. It is a general property of memory that we store only partial descriptions of the things to be remembered, descriptions that are sufficiently precise to work at the time something is learned, but that may not work later on, when new experiences have also been encountered and entered into memory. The descriptions formed to distinguish among the old coins were not precise enough to distinguish between the new one and at least one of the old ones.[5]

Suppose I keep all my notes in a small red notebook. If this is my only notebook, I can describe it simply as my notebook. If I buy several more notebooks, the earlier description will no longer work. Now I must call the first one small or red, or maybe both small and red, whichever allows me to distinguish it from the others. But what if I acquire several small, red notebooks? Now I must find some other means of describing the first book, adding to the richness of the de-

scription and thereby to its ability to discriminate among the several similar items. Descriptions need discriminate only among the choices in front of me, but what works for one purpose may not for another.[6]

THE POWER OF CONSTRAINTS

Back in the good old days of oral tradition (and even today for some cultures), performers traveled around reciting epic poems thousands of lines long. How did they do it? Do some people have huge amounts of knowledge in their heads? Not really. It turns out that external constraints exert powerful control over the permissible choice of words, thus dramatically reducing the memory load.

Consider the constraints of rhyming. If you wish to rhyme one word with another in English, there are usually ten to twenty alternatives. But if you must have a word with a particular meaning to rhyme with another, there are usually no candidates at all. And if there are any, in most cases there is only one. Combining the two constraints of rhyme and meaning can therefore reduce the information about the particular word that must be kept in memory to nothing; as long as the constraints are known, the choice of word can be completely determined. The learning of material like poetry is greatly aided by these kinds of constraints, which work on the general schema for the class of poem, meter, and topic.

Here is an example. I am thinking of three words: one means "a mythical being," the second is "the name of a building material," and the third is "a unit of time." What words do I have in mind? Although you can probably think of three words that fit the descriptions, you are not likely to get the same three that I have in mind. There simply are not enough constraints.

Now try a second task, this time looking for rhyming words. I am thinking of three words: one rhymes with "post," the second with "eel," and the third with "ear." What words am I thinking of?

Suppose I now tell you that the words I seek are the same in both tasks: What is a word that means a mythical being and rhymes with "post"? What word is the name of a building material and rhymes with "eel"? And what word is a unit of time and rhymes with "ear"? Now the task is easy: the joint specification of the words completely constrains the selection.

In the psychology laboratory, people almost never got the correct meanings or rhymes for the first two tasks, but they correctly answered "ghost," "steel," and "year" in the combined task almost always.[7]

The classic study of memory for epic poetry was done by Albert Bates Lord. He went to Yugoslavia and found people who still followed the oral tradition. He demonstrated that the "singer of tales," the person who learns epic poems and goes from village to village reciting them, is really recreating them, composing poetry on the fly in such a way that it obeys the rhythm, theme, story line, structure, and other characteristics of the poem. This is a prodigious feat, but it is not an example of rote memory. Rather, the practice illustrates the immense power of the multiple constraints that allow the singer to listen to another singer tell a lengthy tale once, and then (after a delay of a few hours or a day) apparently recite "the same song, word for word, and line for line."[8] In fact, as Lord points out, the original and new recitations are not the same word for word. But the listener would perceive them as the same, even if the second version were twice as long as the first. They are the same in the ways that matter to the listener: they tell the same story, express the same ideas, and follow the same rhyme and meter. They are the same in all senses that matter to the culture. Lord shows just how the combination of memory for poetics, theme, and style combine with cultural structures into what he calls a formula for producing an appropriate poem, perceived as identical to earlier recitations. The notion that someone should be able to recite word for word is relatively modern. Such a notion can be held only after printed texts become available; otherwise who could judge the accuracy of a recitation? Perhaps more important, who would care? All this is not to detract from the feat. Learning and reciting an epic poem such as Homer's *Odyssey* or *Iliad* is clearly difficult even if the singer is recreating it: there are 27,000 lines of verse in the written version.[9]

Most of us do not learn epic poems. But we do make use of strong constraints that serve to simplify what must be retained in memory. Consider an example from a completely different domain: taking apart and reassembling a mechanical device. Typical items in the home that an adventuresome person might attempt to repair include a door lock, toaster, and washing machine. The device is apt to have tens of parts. What has to be remembered in order to put the parts together again in proper order? Not as much as might appear from an initial analysis. In the extreme case, if there are ten parts, there are 10! (10 factorial: 10

$\times 9 \times 8 \ldots$) different ways in which to reassemble them—a little over 3.5 million alternatives. But never can all possible orderings be produced: there will be a number of physical constraints on the ordering. Some pieces must be assembled before it is even possible to assemble the others. Some pieces are physically constrained from fitting into the spots reserved for others: bolts must fit into holes of an appropriate diameter and depth; nuts and washers must be paired with bolts and screws of appropriate sizes; and washers must always be put on before nuts. There are even cultural constraints: we turn screws clockwise to tighten, counterclockwise to loosen; the heads of screws tend to go on the visible part (front or top) of a piece, bolts on the less visible part (bottom, side, or interior) of a piece; wood screws and machine screws look different and are inserted into different kinds of materials. In the end, the apparently large number of decisions is reduced to only a few choices that should have been learned or otherwise noted during the disassembly. The constraints by themselves are often not sufficient to determine the proper reassembly of the device—mistakes do get made—but the constraints reduce the amount that must be learned to a reasonable quantity.

Memory Is Knowledge in the Head

Remember the story of 'Ali Baba and the forty thieves? 'Ali Baba discovered the secret words that opened the thieves' cave. His brother-in-law, Kasim, forced him to reveal the secret. Kasim then went to the cave.

"When he reached the entrance of the cavern, he pronounced the words, Open Simsim!

"The door immediately opened, and when he was in, closed on him. In examining the cave he was greatly astonished to find much more riches than he had expected from 'Ali Baba's relation. He quickly lade at the door of the cavern as many bags of gold as his ten mules could carry, but his thoughts were now so full of the great riches he should possess, that he could not think of the necessary words to make the door open. Instead of Open Simsim! he said Open Barley! and was much amazed to find that the door remained shut. He named several sorts of grain, but still the door would not open.

"Kasim never expected such an incident, and was so alarmed at the danger he was in that the more he endeavoured to remember the word Simsim the more his memory was confounded, and he had as much forgotten it as if he had never heard it mentioned."

Kasim never got out. The thieves returned, cut off Kasim's head, and quartered his body.[10]

THE CONSPIRACY AGAINST MEMORY

Most of us will not get our heads cut off if we fail to remember a secret code, but it can still be very hard to do. It is one thing to have to memorize one or two secrets: a combination, or a password, or the secret to opening the door. But when the number of secret codes gets too large, memory fails. There seems to be a conspiracy, one calculated to destroy our sanity by overloading our memory. Consider what we are asked to remember in our "convenient" world. A simple search through my own wallet and papers reveals the following things.

- Postal codes ranging in the United States from the "short form" of five digits to the "long form" of nine. Human short-term memory can comfortably retain only a five- to seven-digit number, yet here I am asked to use nine. I need to know the code for where I live, the code for where I work, the codes for my parents and for my children, the codes for my friends, and the codes for anyone with whom I correspond regularly. American codes, such as 92014-6207; British codes, such as WC1N 3BG; Canadian codes, such as M6P2V8. All for the sake of the machinery, and despite the fact that addresses are perfectly sensible and normally unambiguous. But machines have trouble with addresses, whereas they can deal with simple postal codes.
- Telephone numbers, sometimes with area codes and extensions. A seven-digit number becomes ten when the area code is added, and then fourteen when there is a four-digit extension. International codes, with country code and city code, add more digits. How many telephone numbers must I know? More than I wish to contemplate. All my personal contacts. Numbers for information, time, and weather; the special number for emergencies. And I mustn't forget to dial 9 (or, in some cases, 8) so that the call will go outside the institution or company.

- Access numbers for telephone budget cards, so that when I make a long distance call from my university, I can cause the correct account to pay the bill: a five-digit number for each account (and I have four of them). Don't show these to anyone, I am warned. Keep them hidden in a secret place.
- Access numbers for telephone credit cards, so when I travel I can have the bill automatically put on my home telephone number. The codes consist of my home telephone number plus four secret digits. The secret digits aren't even printed on the card: memorize and destroy. But I have six of them (two home phone accounts and four different university phone accounts). If I want to dial a long distance number from a hotel using one of my telephone credit cards, I must dial as many as thirty-six digits.
- Passwords or numbers for bank automatic teller machines, those clever machines that let you put in a card, type in your secret password, and get money. Two bank accounts, two secret passwords. Don't write them down, a thief might see them. Memorize. Memorize.
- Secret passwords for my computer accounts: can't let people steal my valuable data, or perhaps change their course grades, or peek at the examination questions. Make the password at least six characters long, we are told. And no words—words are too easy for someone to discover—make it nonsense. (I cheat and make all my computer accounts use the same password.)
- Driver's license number. When I lived briefly in Texas I couldn't do anything without my driver's license number: not pay for food at the supermarket, not pay the telephone bill, not even open up a bank account. That was one letter, seven digits. Other states have longer numbers.
- Social security numbers for me, my wife, and my children. Nine digits each.
- Passport numbers, again for my whole family.
- My employee number.
- License plate numbers for our cars.
- Birthdays.
- Ages.
- Clothing sizes.
- Addresses.

- Credit card numbers.
- Bah and humbug.

So many of these numbers and codes must be kept secret. Apparently, thieves are everywhere, just waiting for me to write down my secret password or number, anxious to make that phone call on my account or to purchase items with my charge card. There is no way that I can learn all those numbers. And they keep changing, anyway, some of them annually. I even have trouble remembering how old I am: it changes every year too. (Quick: what magic phrase was Kasim trying to remember to open the cavern door?)

How can we remember all these things? Most of us can't, even with the use of mnemonics to make some sense of nonsensical material. Books and courses on improving memory can work, but the methods are laborious to learn and need continued practice to maintain. So we put the memory in the world, writing things down in books, on scraps of paper, even on the backs of our hands. But we disguise them to thwart would-be thieves. That creates another problem: How do we disguise the items, how do we hide them, and how do we remember what the disguise was or where we put them? Ah, the foibles of memory.

Where should you hide something so that nobody else will find it? In unlikely places, right? Money is hidden in the freezer, jewelry in the medicine cabinet or in shoes in the closet. The key to the front door is hidden under the mat or just below the window ledge. The car key is under the bumper. The love letters are in a flower vase. The problem is, there aren't that many unlikely places in the home. You may not remember where the love letters or keys are hidden, but your burglar will. Two psychologists who examined the issue described the problem this way:

"There is often a logic involved in the choice of unlikely places. For example, a friend of ours was required by her insurance company to acquire a safe if she wished to insure her valuable gems. Recognizing that she might forget the combination to the safe, she thought carefully about where to keep the combination. Her solution was to write it in her personal phone directory under the letter S *next to 'Mr. and Mrs. Safe,' as if it were a telephone number. There is a clear logic here: Store numerical information with other numerical information. She was appalled, however, when she heard a reformed burglar on a daytime*

television talk show say that upon encountering a safe, he always headed for the phone directory because many people keep the combination there."[11]

All these numbers to remember add up to unwitting tyranny. It is time for a revolt.

THE STRUCTURE OF MEMORY

"Say aloud the numbers 1, 7, 4, 2, 8. Next, without looking back, repeat them. Try again if you must, perhaps closing your eyes, the better to 'hear' the sound still echoing in mental activity. Have someone read a random sentence to you. What were the words? The memory of the just present is available immediately, clear and complete, without mental effort.

"What did you eat for dinner three days ago? Now the feeling is different. It takes time to recover the answer, which is neither as clear nor as complete a remembrance as that of the just present, and the recovery is likely to require considerable mental effort. Retrieval of the past differs from retrieval of the just present. More effort is required, less clarity results. Indeed, the 'past' need not be so long ago. Without looking back, what were those digits? For some people, this retrieval now takes time and effort."[12]

Psychologists distinguish between two major classes of memory: short-term memory and long-term memory (abbreviated STM and LTM, respectively). The two are quite different. Short-term memory is the memory of the just present. Information is retained in it automatically and retrieved without effort; but the amount of information that can be retained this way is severely limited. Something like five to seven items is the limit of STM, with the number going to ten or twelve if a person also rehearses, mentally repeating the items to be retained. Short-term memory is invaluable in the performance of everyday tasks, in letting us remember words, names, phrases, and parts of tasks. It acts as a working or temporary memory. But the memory is quite fragile. Get distracted by some other activity and, poof, the stuff in STM disappears. It is capable of holding a five-digit postal code or seven-digit telephone number from the time you look them up until the time they are used—as long as no distractions occur. Nine- or

ten-digit numbers give trouble, and when the number starts to exceed that—don't bother. Write it down. Or divide the number into several shorter segments.

Long-term memory is memory for the past. As a rule, it takes time to put stuff away in LTM and time and effort to get it out again. This is how we maintain our experiences, not as an exact recording of the events, but as interpreted through our understanding of them, subject to all the distortions and changes that the human explanatory mechanism imposes upon life. How well we can ever recover experiences and knowledge from LTM is highly dependent upon how the material was interpreted in the first place. What is stored in LTM under one interpretation probably cannot be found later on when sought under some other interpretation. As for how large the memory is, nobody really knows: billions of items, probably. One informed scientist estimates the capacity as a billion (10^9) bits or about 100 million (10^8) items.[13] Whatever the size, it is so large as not to impose any practical limit. The difficulty with LTM is in organization—in getting material in and in figuring out how to retrieve it—not in capacity. Storage and retrieval are easier when the material makes sense, when it fits into what is already known. When the material makes no sense, it will have to be worked on, structured, and interpreted, until finally it can be retained.

Human memory is essentially knowledge in the head, or internal knowledge. If we examine how people use their memories and how they retrieve information, we discover a number of categories. Three are important for us now:

1. *Memory for arbitrary things.* The items to be retained seem arbitrary, with no meaning and no particular relationship to one other or to things already known.
2. *Memory for meaningful relationships.* The items to be retained form meaningful relationships with themselves or with other things already known.
3. *Memory through explanation.* The material does not have to be remembered, but rather can be derived from some explanatory mechanism.

MEMORY FOR ARBITRARY THINGS

Arbitrary knowledge can be classified as the simple remembering of what is to be done, without reliance on an understanding of why or on internal structure. This is how we learned the alphabet and how to tie

a shoelace. It is even how we learned the multiplication tables, that 3 times 2 is 6, although for that we could refer to an external structure. This is how we are expected to learn arbitrary codes to operate the modern, misbegotten telephone system. It is also how we are forced to learn many procedures required of modern technology: "To load this program, put the floppy diskette into drive A and type ALT MODE, CONTROL-SHIFT-X, DELETE." This is rote learning, the bane of modern existence.

Rote learning creates problems. First, because what is being learned is arbitrary, the learning is difficult: it can take considerable time and effort. Second, when a problem arises, the memorized sequence of actions gives no hint of what has gone wrong, no suggestion of what might be done to fix the problem. Although some things are appropriate to learn by rote (the letters of the alphabet, for example), most are not. Alas, it is still the dominant method of instruction in many school systems, and even for much adult training. This is how some people are taught to use computers, or to cook. It is how we have to learn to use some of the new (poorly designed) gadgets of our technology.

Most psychologists would argue that it is not really possible to learn arbitrary associations or sequences. Even where there appears to be no structure, people manufacture some artificial and usually rather unsatisfactory one, which is why the learning is so bad. For our purposes it does not matter whether arbitrary learning is impossible or simply very difficult, the end result is the same: it is not the best way to go, not if there is any choice in the matter. Thus, in teaching the alphabet, we try to make it into a tune, using the natural constraints of rhyme and rhythm to simplify the memory load. People who have learned to use computers or cook by rote are probably not very good. Since they do not understand the reasons for their actions, they must find tasks arbitrary and strange. When something goes wrong, they don't know what to do (unless they've memorized solutions). Although rote learning is at times necessary or efficient—so that emergency procedures for things like high-speed military jet aircraft are handled quickly, automatically when the need arises—on the whole, it is most unsatisfactory.

MEMORY FOR MEANINGFUL RELATIONSHIPS

Most things in the world have a sensible structure, which tremendously simplifies the memory task. When things make sense, they correspond to knowledge that we already have, so the new material can

be understood, interpreted, and integrated with previously acquired material. Now we can use rules and constraints to help understand what things go together. Meaningful structure can organize apparent chaos and arbitrariness.

Remember the discussion of mental models in chapter 2? Part of the power of a good mental model lies in its ability to provide meaning to things. Let's look at an example to show how a meaningful interpretation transforms an apparently arbitrary task into a natural one. Note that the appropriate interpretation may not at first be obvious; it, too, is knowledge and has to be discovered.

A Japanese colleague, call him Mr. Tanaka, had difficulty remembering how to use the turn-signal switch on his motorcycle's left handlebar. Moving the switch forward signaled a right turn, backward a left turn. The meaning of the switch was clear and unambiguous, but the direction in which it should be moved was not. Tanaka kept thinking that because the switch was on the left handlebar, pushing it forward should signal a left turn. That is, he was trying to map the action "push the left switch forward" to the intention "turn left," which was wrong. As a result, he had trouble remembering which switch direction should be used for which turning direction. Most motorcycles have the turn-signal switch mounted differently, rotated 90°, so that moving it left signals a left turn, moving it right a right turn. This mapping is easy to learn (it is a natural mapping). But the turn switch on Tanaka's motorcycle moved forward and back, not left and right. How could he learn it?

Mr. Tanaka solved the problem by reinterpreting the action. Consider the way the handlebars of the motorcycle turn. For a left turn, the left handlebar moves backward. For a right turn, the left handlebar moves forward. The required switch movements exactly paralleled the handlebar movements. If the task is reconceptualized as signaling the direction of motion of the handlebars rather than the direction of the motorcycle, the switch motion can be seen to mimic the desired motion; finally we have a natural mapping. At first, the motion of the switch seemed arbitrary, indirect, and difficult to remember. With the proper interpretation, the switch motion is direct and logical, and, as a result, easy to learn and to use. A meaningful relationship can be indispensable, but you have to have the right one.[14]

Without the proper interpretation, it was difficult to remember the switch directions. With it, both the remembering and the performance of the task became trivial. Note that Tanaka's interpretation of the switch movement did not *explain* anything. It simply let him relate the proper direction to move the switch with the direction in which he was turning the motorcycle. The interpretation is essential, but it should not be confused with understanding.

MEMORY THROUGH EXPLANATION

Now we come to a different, more powerful form of internal memory: understanding. People are explanatory creatures, as I showed in chapter 2. Explanations and interpretations of events are fundamental to human performance, both in understanding the world and in learning and remembering. Here mental models play a major role. Mental models simplify learning, in part because the details of the required behavior can be derived when needed. They can be invaluable in dealing with unexpected situations. Note that the use of mental models to remember (in this case, derive) behavior is not ideal for tasks that must be done rapidly and smoothly. The derivation takes time and requires mental resources, neither of which may be in great supply during critical incidents. Mental models let people derive appropriate behavior for situations that are not remembered (or never before encountered). People probably make up mental models for most of the things they do. This is why designers should provide users with appropriate models: when they are not supplied, people are likely to make up inappropriate ones.[15]

The sewing machine provides a good example of the power of a mental model. A sewing machine is a mysterious beast, managing to loop an upper thread through a lower thread, even though each thread is always connected to its spool or bobbin, respectively. The mental model has to explain how the upper thread goes through the material being sewn, dips under the surface plate, and then loops around the lower thread.

The proper model, it turns out, is something like this. Picture the lower bobbin held gently in the machine by a kind of cup with sloping sides. The cup keeps the bobbin stable, allowing it to rotate so its thread can be unwound. Yet the cup is loose enough so that the upper thread can go inside the cup and around the bobbin—and therefore around the bottom thread. When the upper needle goes through the material and under the plate, a rotating hook grabs its thread and

guides it between the inner walls of the cup and the outer walls of the bobbin case. This helps explain why the machine won't work properly if the bobbin is bent, even if the bobbin still appears to fit and the bottom thread unrolls properly. It explains why dirt on the bobbin or in the cup will mess things up, and why certain kinds of upper thread might cause more trouble than others. (A thick upper thread, especially one that was rough or sticky, might not go smoothly around the bobbin.)

To be honest, I don't know if anything I just said about the failures of bobbins is true. I derived each example from my mental model of a sewing machine. I can't sew. But when Naomi Miyake did her research for her doctoral thesis in my laboratory, she studied people's understanding of sewing and of the machines. The result was twofold: a fine piece of research for her and a mental model for me. So now I can derive what would happen, even if it has never happened to me.

The power of mental models is that they let you figure out what would happen in novel situations. Or, if you are actually doing the task and there is a problem, they let you figure out what is happening. If the model is wrong, you will be wrong too. Am I right about the sewing machine? Decide for yourself: go look at one.

After word got out that I was collecting instances of design peculiarities, a friend reported the following about the sunroof of his new car, an Audi. Supposedly, if the ignition is not on, the sunroof cannot be operated. However, a mechanic explained that you could close the sunroof even without the ignition key if you turned on the headlights and then (1) pulled back on the turn-signal stalk (which normally switches the headlights to high beam), and (2) pushed the close control for the sunroof.

My friend said that it was thoughtful of Audi to provide this override of the ignition key in case the sunroof was open when it started raining. You could close it even if you didn't have your key. But we both wondered why the sequence was so peculiar.

Ever skeptical, I asked to see the manual for the car. The manual was explicit: "You cannot work the sunroof if the ignition is off." A similar statement appeared in the discussion of the electrically powered windows. My friend's mental model was functional: it explained why you would want such a feature, but not how it worked. If the feature was so desirable, why was it not mentioned in the manual?

We searched for another explanation. Perhaps it wasn't a design feature, after all. Perhaps it was an accident of design. Perhaps turning on the lights and pulling back on the stalk connected the electrical power to the car, overriding the fact that the ignition key was off. This would allow the sunroof to work, but only as a by-product of the way the lights were wired.

This model was more specific. It explained what was happening and allowed us to predict that all electrical items should work. So we checked. Turning on the light switch without engaging the ignition did not turn on the headlights; only the parking lights went on. But when we also pulled back on the turn-signal stalk, the headlights did turn on, even though the ignition was off. With the stalk pulled back, the sunroof would close and open. The windows would close and open. The fan on the heating system worked. So did the radio. This was an effective mental model. Now we could understand better what was happening, predict new results, and more easily remember the peculiar set of operations required for the task.

Memory Is Also Knowledge in the World

As we have seen, knowledge in the world, external knowledge, can be very valuable. But it, too, has drawbacks. For one, it is available only if you are there, in the appropriate situation. When you are somewhere else, or if the world has changed meanwhile, the knowledge is gone. The critical memory aids provided by the external information are absent, and so the task or item may not be remembered. A folk saying captures this situation well: *"Out of sight, out of mind."*

REMINDING

One of the most important and interesting aspects of the role of external memory is reminding, a good example of the interplay between knowledge in the head and in the world. Suppose a neighboring family asks you to take them to the airport. You agree to take them next Saturday at 3:30 P.M. Now the knowledge is in your head, but how are you going to remember it at the proper time? You will need to be reminded. There are many strategies for reminding. One is simply to keep the information in your head. If the event is important enough,

you count on having it come repeatedly to mind—what psychologists call rehearsal—so that you can simply assume that there will be no difficulty at all remembering when to leave on Saturday. You can keep the information in your head especially when the event is of great personal importance: suppose you are catching the plane for your first trip to Paris. You won't have any problem remembering. But keeping the knowledge in your head is not ordinarily a good reminding technique.

Suppose the event is not personally important, it is several days away, and you are leading a very busy life. Now you'd better transfer some of the burden of remembering to the outside world. Here is where you use notes to yourself, or pocket and desk calendars or diaries, or electronic alarm clocks that can be set for time of day and date. Or you can ask a friend to remind you. Those of us with secretaries put the burden on them. They, in turn, write notes, enter events on calendars, or set an alarm on the computer system (if it is well enough designed that they can figure out how to work it).

A good reminding method is to put the burden on the thing itself. Do my neighbors want me to take them to the airport? Fine, but they have to call me up the night before and remind me. Do I want to remember to take a book to the university to give to a colleague? I put the book someplace where I cannot fail to see it when I leave the house. A good spot is against the front door of the house. I can't leave without tripping over the book. If I am at a friend's house and I borrow a paper or a book, I remember to take it by putting my car keys on it. Then when I leave, I am reminded. Even if I forget and go out to my car, I can't drive away without the keys.

There are two different aspects to a reminder: the signal and the message. Just as in doing an action we can distinguish between knowing what can be done and knowing how to do it, in reminding we must distinguish between knowing that something is to be remembered and remembering what it is. Most popular reminding devices provide only one of these two critical aspects. The famous "tie a string around your finger" reminder provides only the signal. It gives no hint of what is to be remembered. Writing a note to yourself provides only the message; it doesn't remind you ever to look at it. (Tying a knot in your handkerchief—Carelman's device in figure 3.2—provides neither signal nor message.) The ideal reminder has to have both components: the signal that something is to be remembered, the message of what it is.

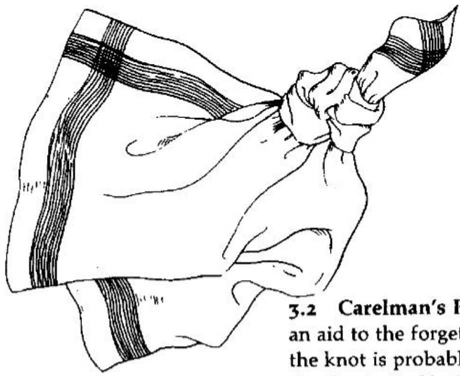

3.2 Carelman's Preknotted Handkerchief. What an aid to the forgetful—except that the act of tying the knot is probably just as useful a memory cue as the knot itself. (Jacques Carelman: "Preknotted Handkerchief" Copyright © 1969–76–80 Jacques Carelman and A. D. A. G. P. Paris. From Jacques Carelman, *Catalog of Unfindable Objects.* Balland, éditeur, Paris-France. Used by permission of the artist.)

The need for timely reminders has created loads of products that make it easier to put the knowledge in the world—alarm clocks, diaries, calendars. A variety of sophisticated watches and small, calculator-sized reminding devices are starting to appear. So far they are limited in power and difficult to use. But I believe there is a need for them. They just need more work, better technology, and better design.

Would you like a pocket-size device that reminded you of each appointment and daily event? I would. I am waiting for the day when portable computers become small enough that I can keep one with me at all times. I will definitely put all my reminding burdens upon it. It has to be small. It has to be convenient to use. And it has to be relatively powerful, at least by today's standards. It has to have a full, standard typewriter keyboard and a reasonably large display. It needs good graphics, because that makes a tremendous difference in usability, and a lot of memory—a huge amount, actually. And it should be easy to hook up to the telephone; I need to connect it to my home and laboratory computers. Of course, it should be relatively inexpensive.

What I ask for is not unreasonable. The technology I need is available today. It's just that the full package has never been put together, partly because the cost in today's world would be prohibitive. But it will exist in imperfect form in five years, possibly in perfect form in ten.

NATURAL MAPPINGS

The arrangement of burners and controls on the kitchen stove provides a good example of the power of natural mappings to reduce the need for information in memory. Without a good mapping, the user cannot readily determine which burner goes with which control. Consider the standard stove with four burners, arranged in the traditional rectangle. If the four controls were truly arbitrary, as in figure 3.3, the user would have to learn each control separately: twenty-four possible arrangements. Why twenty-four? Start with the leftmost control: it could work any of the four burners. That leaves three possibilities for the next leftmost. So there are 12 (4×3) possible arrangements of the first two controls: four for the first, three for the second. The third control could work either of the two remaining burners, and then there is only one burner left for the last control. This makes twenty-four possible mappings between the controls and burners: $4 \times 3 \times 2 \times 1 = 24$. With the completely arbitrary arrangement, the stove is unworkable unless each control is fully labeled to indicate which burner it controls.

Most stoves have controls arranged in a line, even though the burners are arranged rectangularly. Controls are not mapped naturally to burners. As a result, you have to learn which control goes with which burner. Consider how the use of spatial analogies can relieve the memory burden. Start with a partial mapping that is in common use today: the controls are segregated into left and right halves, as in figure 3.4. Now we need know only which left burner each of the two left controls affects and which right burner each right control affects—two alternatives for each of the four burners. The number of possible arrangements is now only four—two possibilities for each side: quite a reduction from the twenty-four. But the controls must still be labeled, which indicates that the mapping is still imperfect. Since some of the information is now in the spatial arrangement, each control need only be labeled back or front; the left and right labels are no longer needed.

What about a proper, full, natural mapping, with the controls spatially arranged in the same pattern as the burners, as in figure 3.5? The organization of the controls now carries all the information required. We know immediately which control goes with which burner. Such is the power of natural mapping. We can see that the number of possible sequences has been reduced from twenty-four to one.[16] If all possible

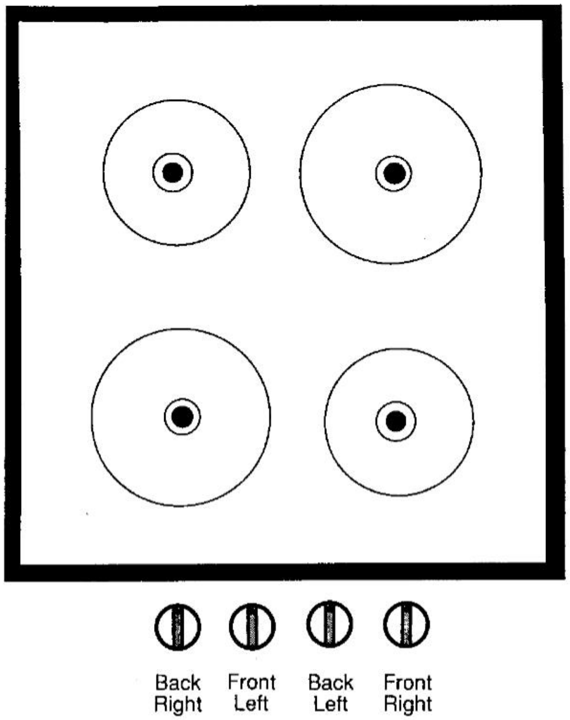

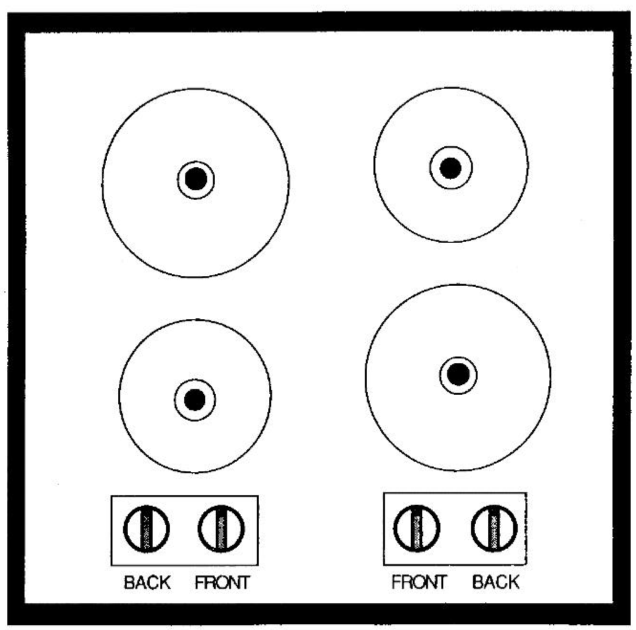

3.3 Arbitrary Arrangement of Stove Controls (top of opposite page). Couple the usual rectangular arrangement of burners with this arbitrary row of controls, and there is trouble: which control goes with which burner? You don't know unless the controls are labeled. The memory load for this arrangement is high: there are twenty-four possible arrangements, and you have to remember which of the twenty-four this one is. Fortunately, the controls are seldom arranged quite this arbitrarily.

3.4 Paired Stove Controls (bottom of opposite page). This is the type of partial mapping of controls to burners in common use today. The two controls on the left work the left burners, and the two controls on the right work the right burners. Now there are only four possible arrangements (two for each side). Even so, confusion is possible (and, I can assure you, it occurs often).

3.5 Full Natural Mapping of Controls and Burners (below). Two of the Possible Ways. There is no ambiguity, no need for learning or remembering, no need for labels. Why can't all stoves be like these?

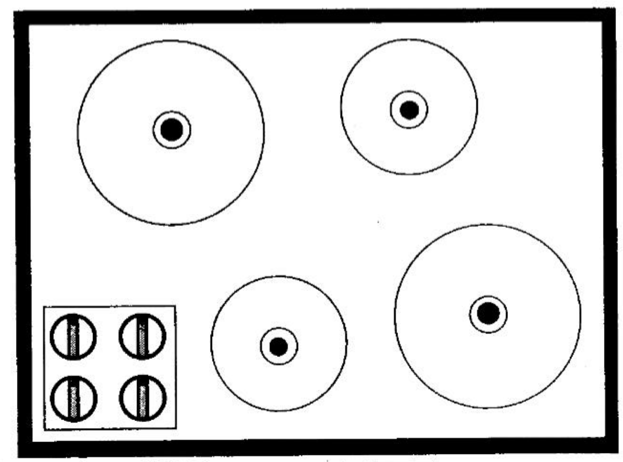

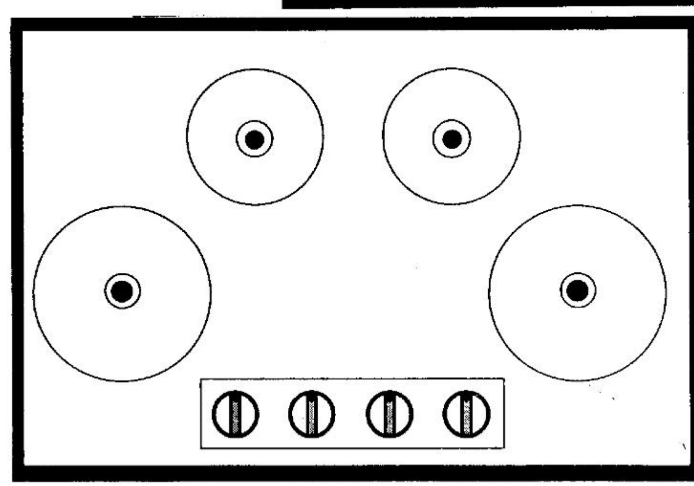

natural mappings were applied in our lives, the cumulative effect would be enormous.

The problem of the stove top may seem trivial, but in fact it is a cause of great frustration for many homeowners. Why do stove designers insist on arranging the burners in a rectangular pattern and the controls in a row? We have known for forty years just how bad such an arrangement is. Sometimes the stove comes with clever little diagrams to indicate which control works which burner. Sometimes there is a short label. But the proper natural mapping requires no diagrams, no labels, and no instructions. There is a simple design principle lurking here:

> If a design depends upon labels, it may be faulty. Labels are important and often necessary, but the appropriate use of natural mappings can minimize the need for them. Wherever labels seem necessary, consider another design.

The shame about stove design is that it isn't hard to do right. Textbooks of ergonomics, human factors, psychology, and industrial engineering all show various sensible solutions. And some stove manufacturers do use good designs. Oddly, some of the very best and the very worst are manufactured by the same companies and are illustrated side by side in the same catalogs.

Why do designers insist on frustrating users? Why do users still purchase stoves that cause so much trouble? Why not revolt and refuse to buy them unless the controls have an intelligent relationship to the burners? I bought a bad one myself.

Usability is not often thought of as a criterion during the purchasing process. Moreover, unless you actually test a number of units in a realistic environment doing typical tasks, you are not likely to notice the ease or difficulty of use. If you just look at something, it appears straightforward enough, and the array of wonderful features seems to be a virtue. You may not realize that you won't be able to figure out how to use those features. I urge you to test products before you buy them. Pretending to cook a meal, or setting the channels on a video set, or attempting to program a VCR will do. Do it right there in the store. Do not be afraid to make mistakes or ask stupid questions. Remember, any problems you have are probably the design's fault, not yours.

A major problem is that often the purchaser is not the user. Appliances may be in a home when people move in. In the office, the purchasing department orders equipment based upon such factors as price, personal relationships with the supplier, and perhaps reliability: usability is seldom considered. Finally, even when the purchaser is the end user, it is sometimes necessary to trade one desirable feature for an undesirable one. In the case of my family's stove, we did not like the arrangement of controls, but we bought the stove anyway: we traded off layout of the burner controls for another feature that was more important to us and available only from one manufacturer. (I return to these issues in chapter 6.)

The Tradeoff between Knowledge in the World and in the Head

Knowledge (or information) in the world and in the head are both essential in our daily functioning. But to some extent we can choose to lean more heavily on one or the other. That choice requires a trade-

3.6 Tradeoffs

PROPERTY	KNOWLEDGE IN THE WORLD	KNOWLEDGE IN THE HEAD
Retrievability	Retrievable whenever visible or audible.	Not readily retrievable. Requires memory search or reminding.
Learning	Learning not required. Interpretation substitutes for learning. How easy it is to interpret information in the world depends upon how well it exploits natural mappings and constraints.	Requires learning, which can be considerable. Learning is made easier if there is meaning of structure to the material (or if there is a good mental model).
Efficiency of use	Tends to be slowed up by the need to find and interpret the external information.	Can be very efficient.
Ease of use at first encounter	High.	Low.
Aesthetics	Can be unaesthetic and inelegant, especially if there is a need to maintain a lot of information. This can lead to clutter. In the end, aesthetic appeal depends upon the skill of the designer.	Nothing need be visible, which gives more freedom to the designer, which in turn can lead to better aesthetics.

off—gaining the advantages of knowledge in the world means losing the advantages of knowledge in the head (figure 3.6).

Knowledge in the world acts as its own reminder. It can help us recover structures that we otherwise would forget. Knowledge in the head is efficient: no search and interpretation of the environment is required. In order to use knowledge in the head we have to get it there, which might require considerable amounts of learning. Knowledge in the world is easier to learn, but often more difficult to use. And it relies heavily upon the continued physical presence of the information; change the environment and the information is changed. Performance relies upon the physical presence of the task environment.

Reminders provide a good example of the relative tradeoffs between the roles of internal versus external knowledge. Knowledge in the world is accessible. It is self-reminding. It is always there, waiting to be seen, waiting to be used. That is why we structure our offices and our places of work so carefully. We put piles of papers where they can be seen, or if we like a clean desk, we put them in standardized locations and teach ourselves (knowledge in the head) to look in these standard places routinely. We use clocks and calendars and notes. Knowledge in the mind is ephemeral: here now, gone later. We can't count on something being present in mind at any particular time, unless it is triggered by some external event or unless we deliberately keep it in mind through constant repetition (which then prevents us from having other conscious thoughts). Out of sight, out of mind.[17]

CHAPTER FOUR

KNOWING WHAT TO DO

"Q. I read a news item about a new videotape-only player and rejoiced when the writer took a healthy swipe at the incomprehensible instructions that accompany VCRs. I can't even set the time of day on mine!

"There are many consumers out here like me—thwarted by an unfathomable machine and baffled by senseless instructions.

"Is there anyone, anywhere who will translate OR give a short course in VCR at play school level?"[1]

Video cassette recorders—VCRs—can be frightening to people who are unfamiliar with them. Indeed, the number of options, buttons, controls, displays, and possible courses of action is formidable. But at least when we have trouble operating a VCR we have something to blame: the machine's bewildering appearance and the lack of clues to suggest what can be done and how to do it. Even more frustrating, however, is that we often have trouble working devices that we expect to be simple.

The difficulty of dealing with novel situations is directly related to the number of possibilities. The user looks at the situation and tries to discover which parts can be operated and what operations can be done.

Problems occur whenever there is more than one possibility. If there is only one part that can be operated and only one possible action to do, there will be no difficulty. Of course, if the designer has been too clever, hiding all the visible clues, the user may believe there are no alternatives and not even know how to begin.

When we encounter a novel object, how can we tell what to do with it? Either we have dealt with something similar in the past and transfer old knowledge to the new object, or we obtain instruction. In these cases, the information we need is in the head. Another approach is to use information in the world, particularly if the design of the new object has presented us with information that can be interpreted.

How can design signal the appropriate actions? To answer the question we build upon the principles discussed in chapter 3. One important set of signals comes through the natural constraints of objects, physical constraints that limit what can be done. Another set of signals comes from the affordances of objects, which convey messages about their possible uses, actions, and functions. A flat plate affords pushing, an empty container affords filling, and so on. Affordances can signal how an object can be moved, what it will support, and whether anything will fit into its crevices, over it, or under it. Where do we grab it, which parts move, and which parts are fixed? Affordances suggest the range of possibilities, constraints limit the number of alternatives. The thoughtful use of affordances and constraints together in design lets a user determine readily the proper course of action, even in a novel situation.

A Classification of Everyday Constraints

To understand the operation of constraints better, I did some simple experiments. I asked people to put things together from the parts given them; they had never seen the finished structure, and they were not even told what they should be constructing.[2] Let me illustrate with one of the examples: building a motorcycle from a Lego set (a children's construction toy).

The Lego motorcycle (figure 4.1) is a simple toy constructed of thirteen parts, some rather specialized. Of the thirteen parts, only two are alike—rectangles with the word *police* on them. One other piece is a blank rectangle of the same size. Three other pieces match in size and shape but are different colors. So there are two sets of three pieces in

4.1 **Lego Motorcycle.** The toy is shown assembled and in pieces. The thirteen parts are so cleverly constructed that even an adult can put them together. The design exploits constraints to specify just which pieces fit where. Physical constraints limit alternative placements. Semantic and cultural constraints provide the necessary clues for further decisions. For example, semantic constraints stop the user from putting the head backward on the body and cultural constraints dictate the placement of the three lights (the small rectangles, which are red, blue, and yellow).

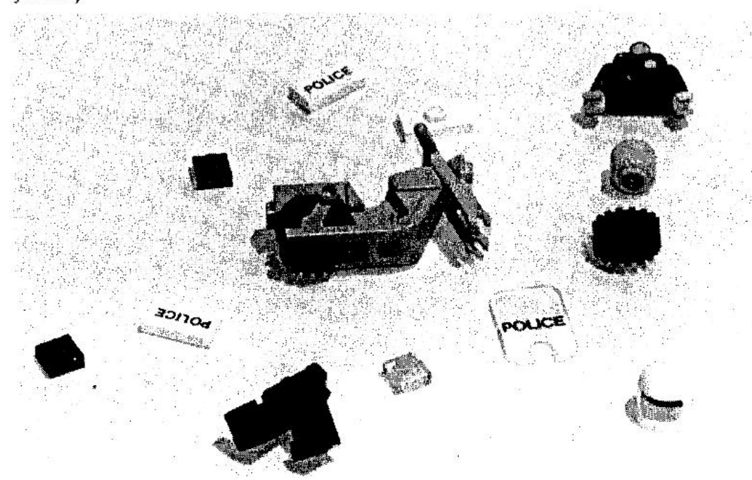

which any of the three pieces are interchangeable, except for the semantic or cultural interpretation of the resulting construction. It turns out that the appropriate role for every single piece of the motorcycle is unambiguously determined by a set of physical, semantic, and cultural constraints. This means that people could construct the motorcycle without any instructions or assistance, although they had never seen it assembled. In this case, construction is entirely natural, if the builder knows about motorcycles and about the cultural assumptions that serve to constrain the placement of parts.

Affordances of the pieces were important in determining just how they fit together. The cylinders and holes characteristic of Lego suggested the major construction rule. The sizes and shapes of the parts suggested their operation. Physical constraints limited what parts would fit together. Other types of constraints also operated; all in all there were four different classes of constraints—physical, semantic, cultural, and logical. These classes are apparently universal, appearing in a wide variety of situations, and sufficient.

PHYSICAL CONSTRAINTS

Physical limitations constrain possible operations. Thus, a large peg cannot fit into a small hole. The motorcycle windshield would fit in only one place, with only one orientation. The value of physical constraints is that they rely upon properties of the physical world for their operation; no special training is necessary. With the proper use of physical constraints there should be only a limited number of possible actions—or, at least, desired actions can be made obvious, usually by being especially salient.

Physical constraints are made more effective and useful if they are easy to see and interpret, for then the set of actions is restricted before anything has been done. Otherwise, the physical constraint prevents the wrong action from succeeding only after it has been tried. The Lego windshield was sometimes tried in the wrong orientation first; the design could have made the correct position more visible. The everyday door key can be inserted into a vertical slot only if the key is held vertically. But this still leaves two possible orientations. A well-designed key will either work in both orientations or provide a clear physical signal for the correct one. Good automobile door keys are made so that orientation doesn't matter. A poorly designed car key can

be yet another of those minor frustrations of everyday life—not so minor, perhaps, when you're standing outside the car in a storm with both arms full of packages.

SEMANTIC CONSTRAINTS

Semantic constraints rely upon the meaning of the situation to control the set of possible actions. In the case of the motorcycle, there is only one meaningful location for the rider, who must sit facing forward. The purpose of the windshield is to protect the rider's face, so it must be in front of the rider. Semantic constraints rely upon our knowledge of the situation and of the world. Such knowledge can be a powerful and important clue.

CULTURAL CONSTRAINTS

Some constraints rely upon accepted cultural conventions, even if they do not affect the physical or semantic operation of the device. One cultural convention is that signs are meant to be read; for the motorcycle, the pieces with the word *police* on them have to be placed right side up. Cultural constraints determine the locations of the three lights, which are otherwise physically interchangeable. Red is the culturally defined standard for a stop light, which is placed in the rear. White or yellow (in Europe) is the standard color for headlights, which go in front. And a police vehicle often has a blue flashing light on top.

Each culture has a set of allowable actions for social situations. Thus, we know how to behave in a restaurant, even one we have never been to before. This is how we manage to cope when our host leaves us alone in that strange room, at that strange party, with those strange people. And this is why we sometimes feel frustrated, so incapable of action, when we are confronted with a restaurant or group of people from an unfamiliar culture, where our normally accepted behavior is clearly inappropriate and frowned upon. Cultural issues are at the root of of many of the problems we have with new machines: there are as yet no accepted conventions or customs for dealing with them.

Those of us who study these things believe that guidelines for cultural behavior are represented in the mind by means of schemas, knowledge structures that contain the general rules and information necessary for interpreting situations and for guiding behavior. In some stereotypical situations (for example, in a restaurant), the schemas may be very specialized. Cognitive scientists Roger Schank and Bob Abelson have proposed that in these cases we follow "scripts" that can guide the sequence of behavior. The sociologist Ervin Goffman calls the social constraints on acceptable behavior frames, and he shows how they govern behavior even when a person is in a novel situation or novel culture. Danger awaits those who deliberately violate the frames for a culture.[3]

Next time you are in an elevator, stand facing the rear. Look at the strangers in the elevator and smile. Or scowl. Or say hello. Or say, "Are you feeling well? You don't look well." Walk up to random passersby and give them some money. Say something like, "You make me feel good, so here is some money." In a bus or streetcar, give your seat to the next athletic-looking teenager you see. The act is especially effective if you are elderly, or pregnant, or disabled.

LOGICAL CONSTRAINTS

In the case of the motorcycle, logic dictated that all the pieces should be used, with no gaps in the final product. The three lights of the Lego motorcycle presented a special problem for many people. They could use the cultural constraint to figure out that the red was the stop light and should go in the rear, that the yellow was the headlight and should go in the front, but what about the blue? Many people had no cultural or semantic information that would help them place the blue light. For them, logic provided the answer: only one piece left, only one possible place to go. The blue light was logically constrained.

Natural mappings work by providing logical constraints. There are no physical or cultural principles here; rather there is a logical relationship between the spatial or functional layout of components and the things that they affect or are affected by. If two switches control two lights, the left switch should work the left light, the right switch the right light. If the lights are mounted one way and the switches another, the natural mapping is destroyed. If two indicators reflect the state of two different parts of a system, the location and operation of the indicators should have a natural relationship to the spatial or functional layout of the system. Alas, natural mappings are not often exploited.

Applying Affordances and Constraints to Everyday Objects

The characteristics of affordances and constraints can be applied to the design of everyday objects, much simplifying our encounters with them. Doors and switches present interesting examples, for poor design causes unnecessary problems for their users. Yet the common problems have simple solutions, which properly exploit affordances and natural constraints.

THE PROBLEM WITH DOORS

In chapter 1 we encountered the sad story of my friend who was trapped between sets of glass doors at a post office, trapped because there were no clues to the doors' operation. When we approach a door, we have to find both the side that opens and the part to be manipulated; in other words, we need to figure out what to do and where to do it. We expect to find some visible signal for the correct operation: a plate, an extension, a hollow, an indentation—something that allows the hand to touch, grasp, turn, or fit into. This tells us where to act. The next step is to figure out how: we must determine what operations are permitted, in part using the affordances, in part guided by constraints.

Doors come in amazing variety. Some open only if a button is pushed, and some don't appear to open at all, having neither buttons, nor hardware, nor any other sign of their operation. The door might be operated with a foot pedal. Or maybe it is voice operated, and we must speak the magic phrase. (*"Open Simsim!"*) In addition, some doors have signs on them: pull, push, slide, lift, ring bell, insert card, type password, smile, rotate, bow, dance, or, perhaps, just ask. Somehow, when a device as simple as a door has to come with an instruction manual—even a one-word manual—then it is a failure, poorly designed.

Appearances deceive. I have seen people trip and fall when they attempted to push open a door that worked automatically, the door opening inward just as they attempted to push against it. On most subway trains, the doors open automatically at each station. Not so in Paris. I watched someone on the Paris Métro try to get off the train and fail. When the train came to his station, he got up and stood patiently *in front of the door, waiting for it to open. It never opened. The train simply started up again and went on to the next station. In the Métro, you have to open the doors yourself by pushing a button, or depressing a lever, or sliding them (depending upon which kind of car you happen to be on).*

Consider the hardware for an unlocked door. It need not have any moving parts: it can be a fixed knob, plate, handle, or groove. Not only will the proper hardware operate the door smoothly, but it will also indicate just how the door is to be operated: it will exhibit the proper affordances. Suppose the door opens by being pushed. The easiest way to indicate this is to have a plate at the spot where the pushing should be done. A plate, if large enough for the hand, clearly and unambiguously marks the proper action. Moreover, the plate constrains the possible actions: there is little else that one can do with a plate except push. Unfortunately, even this simple clue is misused. Doors that should be pulled or slid sometimes have plates (figure 4.2). Doors that should be pushed sometimes have both plates and knobs or a handle and no plate.

The violation of the simple use of constraints on doors can have serious implications. Look at the door in figure 4.3 *A*: this fire exit door has a push bar, a good example of an unambiguous signal to push, and a good design (required by law in the United States) because it forces proper behavior when panicked people press against a door as they attempt to flee a fire. But look again. On which side should you push? There is no way of knowing. Add some paint to the part that is to be pushed, or fasten a plate over it (figure 4.3 *B*): these provide strong cultural signals to guide the action properly. Push bars offer strong physical constraints, simplifying the task of knowing what to do. The use of cultural constraints simplifies the task of figuring out where to do it.

Some hardware cries out to be pulled. Although anything that can be pulled can also be pushed, the proper design will use cultural constraints so that the signal to pull will dominate. But even this can be messed up. I have seen doors with a mixture of signals, one implying push, the other pull. I have watched people passing through the door of figure 4.3 (*A*). And they had trouble, even people who worked in the building and who therefore used the door several times every day.

Sliding doors seem to present special difficulties. In fact, there are several good ways to signal the operation of a sliding door unambiguously. For example, a vertical slit in the door can be used in only one

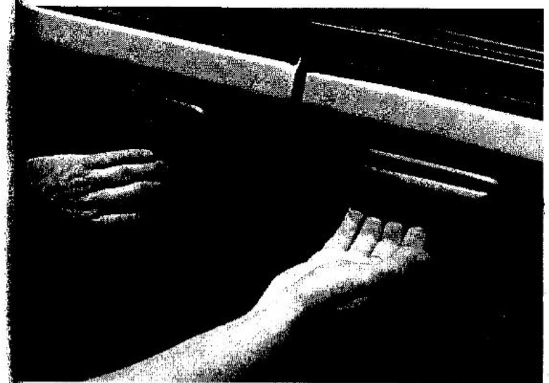

4.2 The Design of Doors. The doors at the left show two excellent examples of design: different handles, side by side on the same automobile, each neatly signaling its proper operation. The vertical placement of the lever on the handle to the left causes the hand to be held in a vertical plane, signifying a slide. The horizontal placement of the lever on the door handle to the right, coupled with the overhang and indentation that neatly afford entrance by the hand, signifies a pull. Two different types of doors, adjacent to one another, and yet there is no confusion between them.

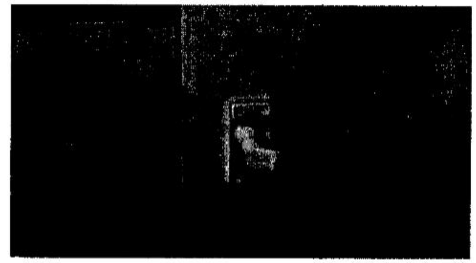

The handle depicted at the left shows inappropriate signals. This form of handle clearly marks grasp, twist, or pull—except that this particular door slides: a classic case of inappropriate design.

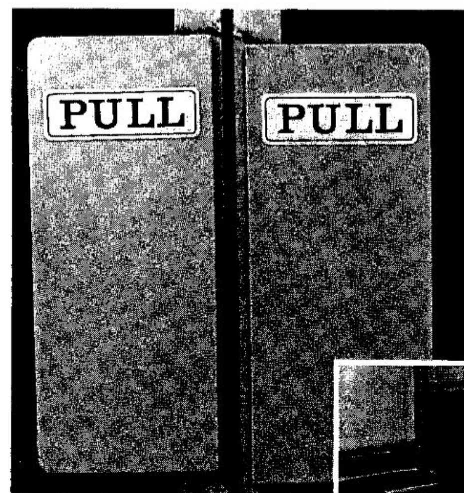

At left and below are photographs of hardware for doors that open by being pulled. The large plates at the left are a signal to push, but in fact the door is supposed to be pulled: no wonder the door needs the signs. The simple U-shaped brackets below is a much better design, but they are ambiguous enough that a sign still seems to be needed. Contrast with the two handles at the top, neither of which needs a sign yet is always operated properly. If a door handle needs a sign, then its design is probably faulty.

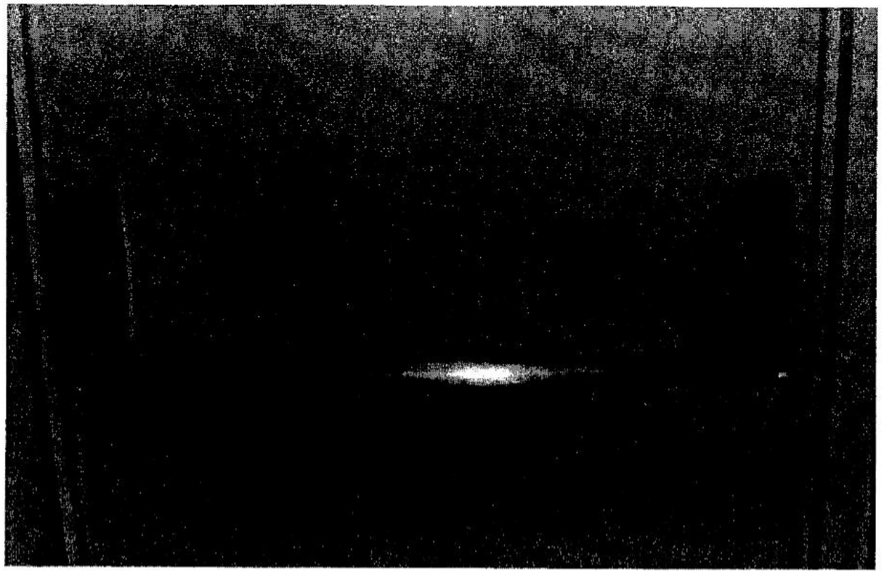

4.3 **Doors in Two Commercial Buildings.** Pushing the bar opens the door, but on which side do you push? Bar *A* (above) hides the signal, making it impossible to know on which side to push. A frustrating door. Bar *B* (below) has a flat plate mounted on the side that is to be pushed; this is a naturally interpreted signal. A nice design, no frustration for the user.

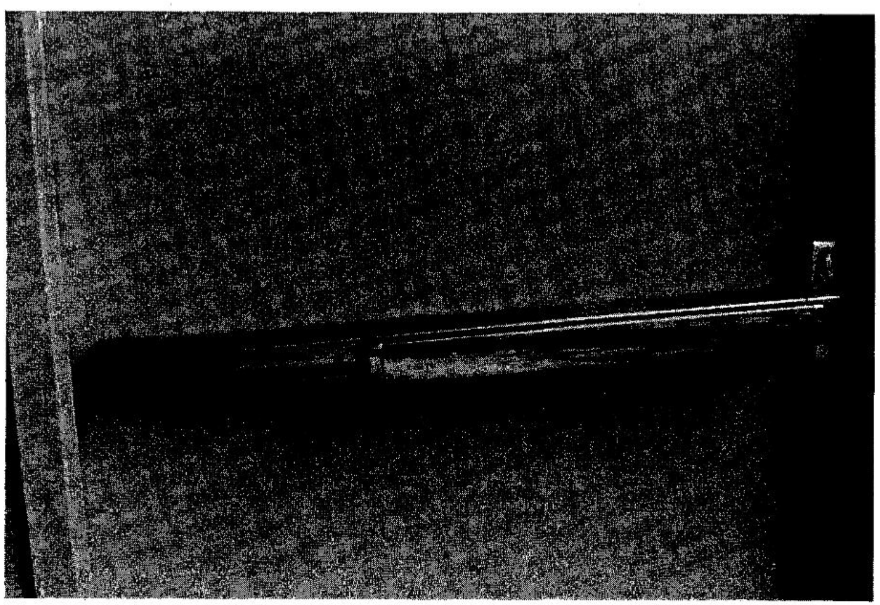

way: the fingers are inserted and the door slid. The location of the slit specifies not only where to exert the force but also in which direction. The critical signal is any depression in the door large enough for the fingers to fit into, but without an overhang. Similarly, any projection will also work, as long as it neither has an overhang nor is appropriate for being grasped with the hand. On a properly designed door, the fingers can exert pressure along the sides of the depression or projection—needed for sliding—but they can't pull or twist. I have seen elegant sliding doors, aesthetically pleasing, yet with clear signals to the user—in a conference room in Italy, on a door on a Métro train in Paris, on some Scandinavian furniture. Yet more often, it seems, sliding doors are built with the wrong signals, with clumsy hardware in positions that jam the fingers. Sliding doors somehow challenge the designer to get them wrong.

Some doors have appropriate hardware, well placed. The outside door handles of most modern automobiles are excellent examples of design. The handles are often recessed receptacles that simultaneously indicate the place and mode of action: the receptacle cannot be used except by inserting the fingers and pulling. Horizontal slits guide the hand into a pulling position; vertical slits signal a sliding motion. Strangely enough, the inside door handles for automobiles tell a different story. Here, the designer has faced a different kind of problem, and the appropriate solution has not yet been found. As a result, although the outside door handles of cars are often excellent, the inside ones are often difficult to find, hard to figure out how to operate, and difficult to use.

Unfortunately, the worst door hardware is found where we spend most of our time: at home and in the office. In many cases, the choice of hardware appears haphazard, used for convenience (or profitability). Architects and interior designers seem to prefer designs that are visually elegant and win prizes. This often means that a door and its hardware are designed to merge with the interior: the door may barely be visible, the hardware merges with door, and the operation is completely obscure. From my experience, the worst offenders are cabinet doors. It is sometimes not even possible to determine where the doors are, let alone whether and from where they are slid, lifted, pushed, or pulled. The focus on aesthetics may blind the designer (and the purchaser) to the lack of usability.

A particularly frustrating design is that of the door that opens outward by being pushed inward. The push releases the catch and ener-

gizes a spring, so that when the hand is taken away the door springs open. It's a very clever design, but most puzzling to the first-time user. A plate would be the appropriate signal, but designers sometimes do not wish to mar the smooth surface of the door. I have such a latch in the glass door of the cabinet in which I store phonograph records. You can see through the door, and it is obvious that there is no room for the door to open inward; to push on the door seems contradictory. New and infrequent users of this door usually reject pushing and open it instead by pulling, which often requires them to use fingernails, knife blades, or more ingenious methods to pry it open.

THE PROBLEM WITH SWITCHES

At any lecture I give, my first demonstration needs no preparation. I can count on the light switches of the room or auditorium to be unmanageable. "Lights please," someone will say. Then fumble, fumble, fumble. Who knows where the switches are and which lights they control? The lights seem to work smoothly only when a technician is hired to sit in a control room somewhere, turning them on and off.

The switch problems in an auditorium are annoying, but similar problems in airplanes and nuclear power plants are dangerous. The controls all look the same. How do the operators avoid the occasional mistake, confusion, or accidental bumping against the wrong control? Or misaim? They don't. Fortunately, airplanes and power plants are pretty robust. A few errors every hour are not important—usually.

One type of popular small airplane has identical-looking switches for flaps and landing gear right next to one another. You might be surprised to learn how many pilots, while on the ground, have decided to raise the flaps and instead raised the wheels. This very expensive error happened frequently enough that the National Transportation Safety Board wrote a report about it. The analysts politely pointed out that the proper design principles to avoid these errors have been known for thirty years. Why were those design errors still being made?

Basic switches and controls should be relatively simple to design well. But there are two fundamental difficulties. The first is the grouping problem, how to determine which switch goes with which function.

The second is the mapping problem. For example, when there are many lights and an array of switches, how can you determine which switch controls which light?

The switch problem becomes serious only where there are many of them. It isn't a problem in situations with one switch, and it is only a minor problem where there are two switches. But the difficulties mount rapidly with more than two switches at the same location. Multiple switches are more likely to occur in offices, auditoriums, and industrial locations than in homes (figure 4.4).

WHICH SWITCH CONTROLS WHICH FUNCTION?

Switches for unrelated functions are often placed together, usually with no distinguishing marks to help the user know which switch controls which function. Designers love rows of identical-looking switches. The switches look good, are easy to mount, are inexpensive to build, and please the aesthetic sensibilities of the viewer. But they

4.4 Typical Audio Mixing Control. This picture was taken in an auditorium in England. Fortunately, errors on panels like these are seldom serious, often not even noted.

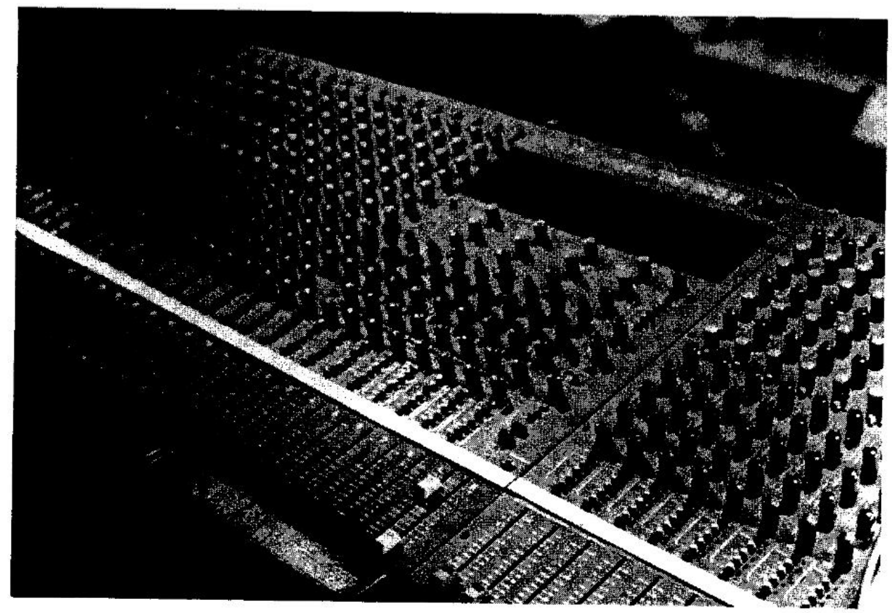

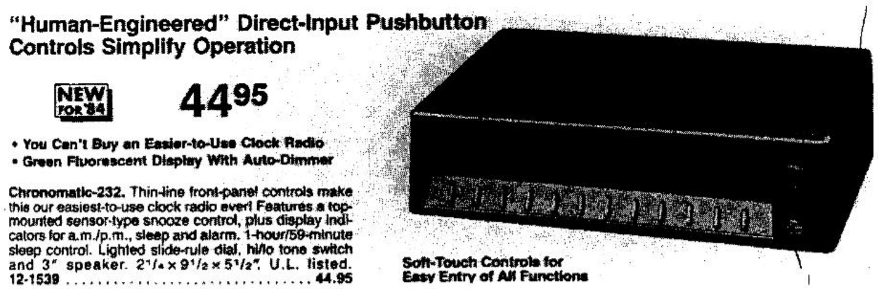

4.5 A Clock Radio, "Human Engineered" to Simplify Operation. Note the row of identical-looking switches. (Copyright Tandy Corporation. Used with permission.)

make it easy to err. With identical switches all in a row, it is difficult to distinguish the switch for the coffee maker from the switch to the central power for the computer. Or the set-the-time switch from the turn-off-the-radio switch (figure 4.5). Or the landing gear switch from the flap control switch.

Consider my car radio: twenty-five controls, many apparently arbitrary. All tiny (so that they will fit the limited space available). Imagine trying to use the radio while driving at high speed, at night. Or in winter when wearing gloves, so that the attempt to push one button succeeds in pushing two, or the attempt to turn the loudness control also adjusts the tone control. You should be able to use things in the dark. A car radio should be usable with a minimum of visual cues. But the radio designers probably designed it in the laboratory, with little or no thought about the car, or the driver. For all I know the design won a prize for its visual aesthetics.

It should go without saying that controls that cause trouble should not be located where they can be operated by accident, especially in the dark, or when the person is trying to use the device without looking. It should go without saying, but in fact, it is necessary to say it.

There is a simple, well-known solution to the grouping problem: set the switches for one set of functions apart from the switches that control other functions. Another solution is to use different types of switches. The solutions can be combined. To solve the problem with the airplane flap and landing gear switches, separate the switches and don't line them up in a row. Also use shape coding: a tire-shaped switch

can control the landing gear, and the flap switch can be a long, thin rectangle—the shape of a flap. Putting controls in different locations makes it less likely that a misaimed hand will throw the wrong switch. And using shape coding means that a potential error may be caught and that the correct switch can be found by feel alone (figure 4.6). That's how to solve this first problem, now let us turn to the other one.

HOW ARE THE SWITCHES ARRANGED?

With the lights in a room, you know that all the switches control lights. But which switch controls which light? Room lights are usually organized in a two-dimensional structure and they are usually horizontal (that is, they are on the ceiling or, if they are lamps, they are placed along the floor or on tables). But switches are usually arranged in a one-dimensional row mounted on the wall, a vertical surface. How can a one-dimensional row of switches map onto a two-dimensional array of lights? And with the switches being mounted on the wall and the

4.6 Make the Controls Look and Feel Different. The control-room operators in a nuclear power plant tried to overcome the problem of similar-looking knobs by placing beer-keg handles over them. This is good design, even if after the fact; the operators should be rewarded. (From Seminara, Gonzales, & Parsons, 1977. Photograph courtesy of Joseph L. Seminara.)

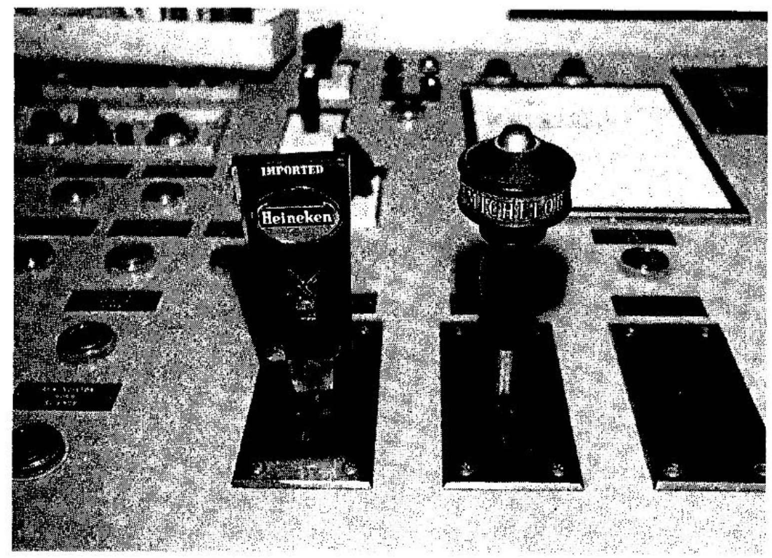

lights being on the ceiling, you have to do a mental rotation of the switches to get them to conform to the lights. The mapping problem is unsolvable with the current structure of switches.

Electricians usually try to lay out the switches in the same order as the lights they control, but the mismatch in the spatial arrangement of the lights and the switches makes it difficult, if not impossible, to produce a full natural mapping. Electricians have to use standard components, and the designers and manufacturers of those standard components worried only about fitting the proper number of switches into them safely. Nobody thought about how the lights were to be arranged or how the switches ought to be laid out.

My house was designed by two brash young architects, award winning, who, among other things, liked neat rows of light switches. We got a horizontal row of four identical switches in the front hall, a vertical column of six identical switches in the living room. "You will get used to it," the architects assured us when we complained. We never did. Finally we had to change the switches, making each one different. Even so we made lots of mistakes.

In my psychology laboratory, the lights and their switches were located in many different places, yet most people wanted to control the lights upon entering the area. The area is large, with three major hallways and approximately fifteen rooms. Moreover, this floor of the building has no windows, so it is dark unless the lights are turned on.

If light switches are placed on the wall, there is no way they can exactly correspond in position to the placement of the lights. Why place the switches flat against the wall? Why not redo things? Why not place the switches horizontally, in exact analogy to the things being controlled, with a two-dimensional layout so that the switches can be placed on a floorplan of the building in exact correspondence to the areas that they control? Match the layout of the lights with the layout of the switches: the principle of natural mapping. In my laboratory, as in my home, the solution was to construct a simple switchplate that mirrored the physical arrangement of the area, with small light switches placed in relevant locations.[4] Figure 4.7 shows the situation at my home, and figure 4.8 shows what we did at the laboratory.

How well do the new switch arrangements work? Quite well, I am happy to report. One laboratory user sent me the following note:

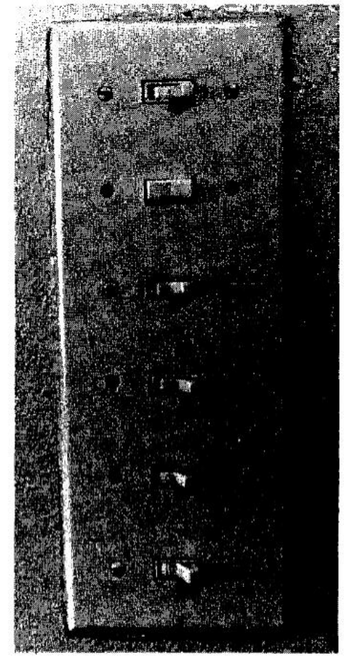

4.7 The vertical array of six switches at the right is what our architects provided to control the lights in our odd-shaped living room. We could never remember which switch did what.

The photograph below shows our solution: switches arranged to match the room layout. (One more switch, for a projection screen, will be mounted on the vertical plate just above the light switches. The switch panel was constructed for the author by David Wargo.)

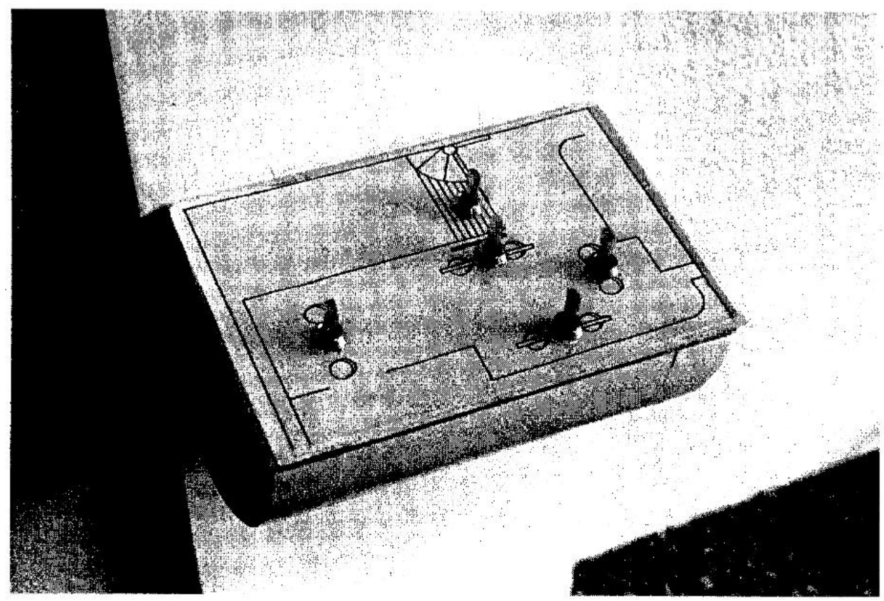

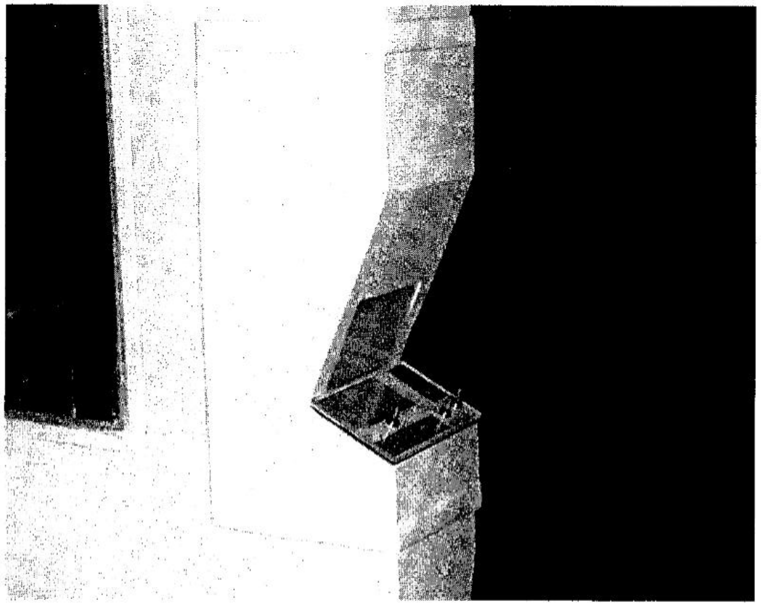

4.8 The original layout of switches in my laboratory had the light switches scattered. We put all the switches in one convenient location, arranged on a floor plan of the laboratory. (The switch panel was constructed by David Wargo.)

"You know, I actually kind of like those new switches now—they seem easy to use, and it's nice to have all the switches in one location when you first walk in. You can just sort of swipe at them on your way past and light up the area you want—very quick. So while I was worried they wouldn't be advantageous for the experienced user, I was wrong."

Can the new switches be used everywhere? Probably not. But there is no reason they couldn't be widely adopted. There are a series of technical problems still to be addressed: builders and electricians need standardized components. How about making up standard light switch boxes, made to be mounted *on* the wall (instead of *in* the wall as they are today), where the switches are mounted on the top of the box, on the horizontal surface. And on the top, make up a matrix of supports so that there can be free, relatively unrestricted placement of the switches in whatever pattern best suits the room. Use smaller switches

if necessary. Maybe get rid of those standardized light plates. The matrix design would require drilling holes differently for each room, but if the switches were designed to fit into standard sized circular or rectangular holes, the holes could be drilled or punched quite easily.

My suggestion requires that the switch box stick out from the wall, whereas today's boxes are mounted so that the switches are flush with the wall. Some might consider my solution ugly. Well, then, indent the boxes, placing them *in* the wall. After all, if there is room inside the wall for the existing switch boxes, there is also room for an indented horizontal surface. Or mount the switches on a little pedestal, or on a ledge.

Visibility and Feedback

So far we have concentrated upon constraints and mappings. But for knowing what to do there are other relevant principles, too, especially visibility and feedback:

1. *Visibility.* Make relevant parts visible.
2. *Feedback.* Give each action an immediate and obvious effect.

When we use a novel object, a number of questions guide our actions:

- Which parts move; which are fixed?
- Where should the object be grasped? What part is to be manipulated? What is to be held? Where is the hand to be inserted? If it is speech sensitive, where does one talk?
- What kind of movement is possible: pushing, pulling, turning, rotating, touching, stroking?
- What are the relevant physical characteristics of the movements? With how great a force must the object be manipulated? How far can it be expected to move? How can success be gauged?
- What parts of the object are supporting surfaces? How much size and weight will the object support?

The same kinds of questions arise whether we are trying to decide what to do or attempting to evaluate the results of an action. In examining the object, we have to decide which parts signify the state of the object and which are solely decorative, or nonfunctional, or part of the background or supports. What things change? What has changed over the previous state? Where should we be watching or listening to detect any changes? The important things to watch should be visible and clearly marked; the results of any action should be immediately apparent.

MAKING VISIBLE THE INVISIBLE

The principle of visibility is violated over and over again in everyday things. In numerous designs crucial parts are carefully hidden away. Handles on cabinets distract from some design aesthetics, and so they are deliberately made invisible or left out. The cracks that signify the existence of a door can also distract from the pure lines of the design, so these significant cues are also minimized or eliminated. The result can be a smooth expanse of gleaming material, with no sign of doors or drawers, let alone of how those doors and drawers might be operated. Electric switches are often hidden: many electric typewriters have the on/off switch hidden underneath; many computers and computer terminals have the on/off switch in the rear, difficult to find and awkward to use;[5] and the switches that control kitchen garbage disposal units are often hidden away, sometimes nearly impossible to find.

Many systems are vastly improved by the act of making visible what was invisible before. Consider the VCR.

> *"UMPTEEN-DAY-UMPTEEN-EVENT PROGRAMMING. Because time-shifting is so popular, manufacturers and retailers play up a VCR's ability to record automatically. The typical VCR can record four events (video jargon for programs) over a 4-day span. . . .*
>
> *"It's one thing to know that a VCR can record eight events in 14 days. It's quite another to make the machine behave. You have to go through a tedious series of steps to tell the VCR when to start recording, what channel to record, how long to run the tape, and so on.*
>
> *"Some VCR's are much easier to program than others. . . . Best of all, we think, is a feature called on-screen programming. Commands that appear on the TV screen help you enter the time, date, and channel of the program you want to tape."*[6]

As the quotation from *Consumer Reports* indicates, the act of setting up these units to do the recording is horribly complex and difficult. The same article later warns that if you are not careful in your selection, "you could wind up with a VCR that brings out fear and loathing whenever you try to change the channel resets or set it up to record a program when you are away." It does not take much examination to discover the reason for the difficulties: there is no visual feedback. As a result, users (1) have trouble remembering their place in the lengthy sequence of required steps; (2) have trouble remembering what next needs to be done; and (3) cannot easily check the information just entered to see if it is what was intended, and then cannot easily change it, if they decide it is wrong.

The gulfs both in execution (the first two problems) and in evaluation (the last problem) are significant for these VCRs. Both can be bridged by the use of a display. Displays often cost money and take up room, which is why designers hesitate to use them, but in the case of a VCR, a display device is usually already available: the TV set. And, indeed, those VCRs that can be programmed through the use of an on-screen TV display are much easier to use. Visibility makes all the difference.

NOTHING SUCCEEDS LIKE A GOOD DISPLAY

Over and over again we find unwarranted complexity that could be avoided were the device to contain a good display. With the modern telephone (see chapter 1), a display that could prompt the user through the series of steps required for programming would make the difference between a valuable, usable system and a next-to-useless one. So, too, with any device of complexity, whether it be the washing machine, microwave oven, or office copying machine. Nothing succeeds like visual feedback, which in turn requires a good visual display.

WHAT CAN BE DONE?

New technologies, especially the inexpensive microprocessors available today (the heart of the computer) make possible the incorporation of powerful and intelligent systems even in simple, everyday things, from toys to kitchen appliances to office machines. But new capabilities must be accompanied by appropriate displays, also now relatively inexpensive. I asked the students in one of my classes to generate some possibilities for adding visibility to everyday devices. Here are some of them:

- *Display the song titles for compact discs.* Why not take advantage of the storage capacity of an audio compact disc (CD) and have it display not only the number of the song or track (as it now does) but also the title? Each title could be accompanied by other information, such as performers, composer, or playing time. Thus, in programming the CD, you could select by name rather than by number, and you would always know what you were hearing.

- *Display the names of television programs.* If each television station would also broadcast its station identification and the title of the current program, the viewer who tuned in during the middle of a show could easily find out what it was. The information could be sent in computer-readable format during the retrace interval (the time that the beam is off the screen).
- *Print the cooking information for foods on the food package in computer-readable form.* This is a scheme for bypassing the need to make things visible. The cooking of frozen foods often requires several different cooking times, waiting times, and heat settings. The programming is complex. If the cooking information were on the package in machine-readable form, one could put the food in the microwave oven, pass a scanner over the printed information, and let the oven program itself.

USING SOUND FOR VISIBILITY

Sometimes things can't be made visible. Enter sound: sound can provide information available in no other way. Sound can tell us that things are working properly or that they need maintenance or repair. It can even save us from accidents. Consider the information provided by:

- The click when the bolt on a door slides home
- The "zzz" sound when a zipper works properly
- The "tinny" sound when a door doesn't shut right
- The roaring sound when a car muffler gets a hole
- The rattle when things aren't secured
- The whistle of a tea kettle when the water boils
- The click when the toast pops up
- The increase in pitch when a vacuum cleaner gets clogged
- The indescribable change in sound when a complex piece of machinery starts to have problems

Many devices do use sound, but only for signals. Simple sounds, such as buzzers, bells, or tones. Computers use bleeping, whining, and

clicking sounds. This use of sound is valuable and serves an important function, but it is very limited in power; it is as if the use of visual cues were limited to different colored, flashing lights. We could use sound for much more communciation than we do.

These days computers produce several sounds, and keypads, microwave ovens, and telephones beep and burp. These are not naturalistic sounds; they do not convey hidden information. When used properly, a beep can assure you that you've pressed a button, but the sound is as annoying as informative. Sounds should be generated so as to give information about the source. They should convey something about the actions that are taking place, actions that matter to the user but that would otherwise not be visible. The buzzes, clicks, and hums that you hear while a telephone call is being completed are one good example: take out those noises and you are less certain that the connection is being made.

Bill Gaver, who has been studying use of sound in my laboratory, points out that real, natural sound is as essential as visual information because sound tells us about things we can't see, and it does so while our eyes are occupied elsewhere. Natural sounds reflect the complex interaction of natural objects: the way one part moves against another; the material of which the parts are made—hollow or solid, metal or wood, soft or hard, rough or smooth. Sounds are generated when materials interact, and the sound tells us whether they are hitting, sliding, breaking, tearing, crumbling, or bouncing. Moreover, sounds differ according to the characteristics of the objects, according to their size, solidity, mass, tension, and material. And they differ with how fast things are going and how far away from us they are.

If they are to be useful, sounds must be generated intelligently, with an understanding of the natural relationship between the sounds and the information to be conveyed. Sounds on artificial devices should be as useful as sounds in the real world. Gaver has proposed that sound could play an important role in computer-based applications. Here, rich, naturalistic sounds could serve as auditory icons, caricatures of naturally occurring sounds that could provide information about the concepts being represented not easily conveyed in other ways.[7]

You have to be very careful with sound, however. It easily becomes cute rather than useful. It can annoy and distract as easily as it can aid. One of the virtues of sounds is that they can be detected even when attention is applied elsewhere. But this virtue is also a deficit, for sounds are often intrusive. Sounds are difficult to keep private unless the intensity is low or earphones are used. This means both that neighbors may be annoyed and that others can monitor your activities. The use of sound to convey information is a powerful and important idea, but still in its infancy.

Just as the presence of sound can serve a useful role in providing feedback about events, the absence of sound can lead to the same kinds of difficulties we have already encountered from a lack of feedback. The absence of sound can mean an absence of information, and if feedback from an action is expected to come from sound, silence can lead to problems.

I once stayed in the guest apartment of a technological institute in the Netherlands. The building was newly completed, with many interesting architectural features. The architect had gone to great lengths to keep the noise level low; the ventilation system could not be heard. In similar fashion, the ventilation for the room came and went through invisible slots in the ceiling (so I am told; I never did find them).

All was fine until I took a shower. The bathroom seemed to have no ventilation at all, so everything became wet, then eventually cold and clammy. There was a switch in the bathroom that I thought might be the control for an exhaust fan. When I pushed the switch, a light on it came on and stayed on. Further pushing had no effect.

I noticed that whenever I returned to the apartment after an absence, the light would be off. So each time I entered the apartment, I went into the bathroom and pushed the button. By listening closely, I could hear a slight "thump" in the distance the first time the button was depressed. I decided it was some kind of signal. Perhaps it was a call button, summoning the maid, or the janitor, or maybe even the fire department (though no one showed up). I did also consider that it might control a ventilation system, but I could hear no flow of air. I examined the inside of the entire bathroom with care, trying to find an air inlet. I even got a chair and a flashlight and examined the ceiling. Nothing.

At the end of my stay, the person driving me to the airport, explained that the button controlled the exhaust fan. The fan was on as long as the light was on, and it turned off, automatically, in about five minutes. The architect was very good at disguising the ventilation system and at keeping the noise level down.

Here is a case where the architect was too successful: the feedback was clearly lacking. The light was not enough—in fact, it was quite misleading. Noise would have been welcome. It would have signaled that there really was ventilation.