7
Being Analog

We are analog beings trapped in a digital world, and the worst part is, we did it to ourselves.

We humans are biological animals. We have evolved over millions of years to function well in the environment, to survive. We are analog devices following biological modes of operation. We are compliant, flexible, tolerant. Yet we have constructed a world of machines that requires us to be rigid, fixed, intolerant. We have devised a technology that requires considerable care and attention, that demands to be treated on its own terms, not ours. We live in a technology-centered world where the technology is not appropriate for people. No wonder we have such difficulties.

Here we are, wandering about the world, bumping into things, forgetful of details, with a poor sense of time, a poor memory for facts and figures, unable to focus attention on a topic for more than a short duration, reasoning by example rather than by logic, and drawing upon our admittedly deficient memories of prior experience. When viewed this way, we seem rather pitiful. No wonder we have constructed a set of artificial devices that are very much not in our own image. We have constructed a world of machinery in which accuracy and precision matter. Time matters. Names, dates, facts, and figures matter. Accurate memory matters. Details matter.

Figure 7.1
Treating people like machines. The foreign service section of switchboard in the American Telephone and Telegraph Company, New York City, 1929. (Photograph courtesy of Corbis-Bettmann.)

All the things we are bad at matter, all the things we are good at are ignored. Bizarre.

Making Sense of the World

People excel at perception, at creativity, at the ability to go beyond the information given, making sense of otherwise chaotic events. We often have to interpret events far beyond the information available, and our ability to do this efficiently and effortlessly, usually without even being aware that we are doing so, greatly adds to our ability to function. This ability to put together a sensible, coherent image of the world in the face of limited evidence allows us to anticipate and predict events, the better to cope with an ambiguous, ever-changing world.

Here's a simple test of your memory:

How many animals of each type did Moses take on the ark?

What's the answer? How many animals? Two? Be careful: What about an amoeba, a sexless, single-celled animal that reproduces by dividing itself into two cells? Did he need to take two of these?

Answer: None. No animals at all. Moses didn't take any animals onto the ark. It was Noah.

Some of you were fooled. Why? Because people often hear what is intended, not what is said. In normal language, people ask real questions that have real answers and real meaning. It is only psychology professors and jokesters who ask trick questions. If you spotted the trick, it is because you were unnaturally suspicious or alert. We don't need such alertness in normal human interaction. Those of you who were fooled responded normally: That is how we are meant to be.

Your mind interpreted the question meaningfully, making sense of the information. It may have confused "Moses" with "Noah," but it was aided by the fact that those names have a lot of similarity: both are short, with two syllables. Both are biblical, from the Old Testament. In normal circumstances, the confusion would be beneficial, for it is the

sort of error that a speaker might make, and it is useful when a listener can go beyond superficial errors.

Note that the ability to be insensitive to simple speech errors does not mean that people are readily fooled. Thus, you would not have been fooled had I asked:

How many animals of each type did Clinton take on the ark?

The name Clinton is not sufficiently close to the target: it requires a biblical name to fool you.[1] From a practical point of view, although a speaker might say "Moses" when "Noah" was intended, it is far less likely that someone would mistakenly say a nonbiblical name such as "Clinton." The automatic inaccurate interpretation of the original question is intelligent and sensible. The fact that the first question can fool people is a testament to our powers, not an indictment of them. Once again, in normal life, such corrections are beneficial. Normal life does not deliberately try to fool us. Take note of this example, for it is fundamental to understanding people and, more important, to understanding why computers are so different from people, why people and today's technology are such a bad match.

Why do accuracy and precision matter? In our natural world, they don't. We are approximate beings; we get at the meanings of things, and for this, the details don't much matter. Accurate times and dates matter only because we have created a culture in which these things are important. Accurate and precise measurements matter because the machines and procedures we have created are rigid, inflexible, and fixed in their ways, so if a measurement is off by some tiny fraction, the result can be a failure to operate. Worse yet, it can cause a tragic accident.

People are compliant: We adapt ourselves to the situation. We are flexible enough to allow our bodies and our actions to fit the circumstances. Animals don't require precise measurements and high accuracy to function. Machines do.

The same story is true of time, of facts and figures, and of accurate memory. These matter only because the mechanical, industrialized society created by people doesn't match people. In part, this is because we don't know how to do any better. Can we build machines that are as compliant and flexible as people? Not today. Biology doesn't build; it grows, it evolves. It constructs life out of flexible parts. Parts that are self-repairable. We don't know how to do this with our machines. We build information devices only out of binary logic, with its insistence upon logic and precision. We invented the artificial mathematics of logic the better to enhance our own thought processes.

The dilemma facing us is the horrible mismatch between the requirements of these human-built machines and human capabilities. Machines are mechanical, we are biological. Machines are rigid and require great precision and accuracy of control. We are compliant. We tolerate and produce huge amounts of ambiguity and uncertainty, very little precision and accuracy. The latest inventions of humankind are those of the digital technology of information processing and communication, yet we ourselves are analog devices. Analog and biological.

An analog device is one in which the representation of information corresponds to its physical structure. In an analog recording the stored signal varies in value precisely in the same way as sound energy varies in time. A phonograph recording is analog; it works by recreating the variations in sound energy by wiggles and changes of depth in the groove. In a tape recording, the strength of the magnetic field on the tape varies in analogous fashion to the sound energy variations. These are analog signals.

Digital signals are entirely different. Here, what is recorded is an abstraction of the real signals. Digital encoding was invented mainly to get rid of noise. In the beginning, electrical circuits were all analog. But electrical circuits are noisy, meaning they are susceptible to unwanted voltage variations. The noise gets in the way, mostly because the circuits are unable to distinguish between the stuff that matters and the stuff that doesn't.

Enter the digital world. Instead of using a signal that is analogous to the physical event, the event is transformed into a series of numbers that describes the original. In high-quality recording of music, the

sound energy is sampled over 40,000 times each second, transformed into numbers that represent the energy value at the time each sample was made. The numbers are usually represented in the form of binary digits rather than the familiar decimal ones, which means that any digit can have only one of two states, 0 or 1, rather than the ten possible states of a decimal digit. When there are only two states to be distinguished between, the operation is far simpler and less subject to error than when it has to determine a precise value, as is required with an analog signal. Binary signals are relatively insensitive to noise.

As you can imagine, to record and playback a digital representation of sound waves requires a lot of processing. It is necessary to transform the sound into numbers, store the numerical digits, and then retrieve and restore them back to sound energy. Such rapid transformation wasn't possible at an affordable price until recently, which is why the emphasis on digital signals seems new. It is only recently that the technology was capable of high-quality digital encoding of audio and television signals, although the concept is old.

There are a number of common misconceptions about digital and analog signals. One is that *analog* means continuous, whereas *digital* means discrete. Although this is often the case, it is not the basis for the distinction. Think of *analog* as meaning *analogous:* analogous to the real world. If the real world event is discrete, so too will be the analog one. If the physical process is continuous, then so too will be the analog one. Digital, however, is always discrete: one of a limited number of values, usually one of two, but occasionally one of three, four, or ten.

A widespread misconception is that digital is somehow good, analog bad. This just isn't so. Yes, digital is good for our contemporary machines, but analog might be better for future machines. And analog is certainly far better for people. Why? Mainly because of the impact of noise.

We have evolved to match the world. If you want to understand how human perception works, it helps to start off by understanding how the world of light and sound works, because the eyes and ears have evolved to fit the nature of these physical signals. What this means is that we interact best with systems that are either part of the real world or analogous to them—analog signals.

Analog signals behave in ways human beings can understand. A slight error or noise transforms the signals in known ways, ways the body has evolved to interpret and cope with. If there is some noise in a conventional television signal, encoded in analogical form, we see some noise on the screen. Usually we can tolerate the resulting image, at least as long as we can make sense of it. Small amounts of noise have slight impact. People are analog, able to extract meanings despite noise and error. As long as the meanings are unchanged, the details of the signals do not matter. They are not noticed, they are not remembered.

In a digital signal, the representation is so arbitrary that a simple error can have unexpected consequences. Digital encodings use compression technologies that eliminate redundancy. Digital television signals are compressed to save space and bandwidth, the most common scheme being the algorithms devised by the Motion Picture Expert Group (MPEG). If any information is lost, it takes a while before the system resends enough information to allow recovery. MPEG encoding breaks up the picture into rectangular regions. Noise can make it impossible for the system to reconstruct an entire region. As a result, when the image is noisy, whole regions of the screen break up and distort in ways the human brain cannot reconstruct, and it takes a few seconds until the picture reforms itself.

The real problem with being digital is that it implies a kind of slavery to accuracy, a requirement that is most unlike the natural workings of the person. It is perfectly proper and reasonable for machines to use digital encodings for their internal workings. Machines do better with digital encoding. The problem comes about in the form of interaction between people and machines. People do best with signals and information that fit the way they perceive and think, which means analogous to the real world. Machines do best with signals and information that are suited for the way they function, which means digital, rigid, precise. So when the two have to meet, which side should dominate? In the past, it has been the machine that dominates. In the future, it should be the person. Stay tuned for chapter 9.

Humans versus Computers[2]

The ever-increasing complexity of everyday life brings with it both great opportunities and major challenges. One of the challenges, that the brain does not work at all like a computer, also provides us with an opportunity: the possibility of new modes of interaction that allow us to take advantage of the complementary talents of humans and machines.

The modern era of information technology has been with us but a short time. Computers are less than a century old. The technology has been constructed deliberately to produce mechanical systems that operate reliably, algorithmically, and consistently. They are based upon mathematics, or more precisely, arithmetic in the case of the first computing devices and logic in the case of the more modern devices.

Contrast this with the human brain. Human beings are the result of millions of years of evolution, where the guiding principle was survival of the species, not efficient, algorithmic computation. Robustness in the face of unexpected circumstances plays a major role in the evolutionary process. Human intelligence has coevolved with social interaction, cooperation, rivalry, and communication. Interestingly enough, the ability to deceive seems to have been one driving force. Only the most intelligent of animals is able to employ a sophisticated level of intentional, purposeful deception. Only the most sophisticated animal is capable of seeing through the deceit. Sure, nature also practices deception through camouflage and mimicry, but this isn't willful and intentional. Primates are the most skilled at intentional, willful deception, and the most sophisticated primate—the human—is the most sophisticated deceiver of all.

Note that some deception is essential for the smooth pursuit of social interaction: the "white lie" smoothes over many otherwise discomforting social clashes. It is not always best to tell the truth when people ask how we like their appearance, or their presentation, or the gift they have just given us. One could argue that computers won't be truly intelligent or social until they, too, are able to deceive.

We humans have learned to control the environment. We are the masters of artifacts. Physical artifacts make us stronger, faster, and more comfortable. Cognitive artifacts make us smarter. Among cognitive artifacts are the invention of writing and other notational systems, such as those used in mathematics, dance, and musical transcription. The result of these inventions is that our knowledge is now cumulative; each generation grows upon the heritage left behind by previous generations. This is the good news. The bad news is that the amount to be learned about the history, culture, and techniques of modern life increases with time. It now takes several decades to become a truly well educated citizen. How much time will be required in fifty years? In one hundred years?

The biological nature of human computation, coupled with the evolutionary process by which the brain has emerged, leads to a very different style of computation from the precise, logic-driven systems that characterize current computers. The differences are dramatic. Computers are constructed from a large number of fast, simple devices, each following binary logic and working reliably and consistently. Errors in the operation of any of the underlying components are not tolerated, and they are avoided either by careful design to minimize failure rates or through error-correcting coding in critical areas. The remarkable power of the computer is a result of the high speed of relatively simple computing devices.

Biological computation is performed by a very large number of slow, complex devices—neurons—each doing considerable computation and operating through electrochemical interactions. The power of the computation is a result of the highly parallel nature of the computation and the complex computations done by each of the billions of neural cells. Moreover, the cells are bathed in fluids whose chemistry can change rapidly, providing a means for rapid dispersion of hormones and other signals to the entire system, chemicals that are site-specific. Think of it as a packet-switching deployment of chemical agents. The result is that the computational basis is dynamic, capable of rapid, fundamental change. Affect, emotion, and mood all play a powerful—and as yet poorly understood—role in human cognition. Certainly all of us have experienced the tension when logic dictates one course of action but

mood or emotion another. More often than not, we follow mood or emotion.

Whatever the mode of computation—and the full story is not yet known—it is certainly not binary logic. Each individual biological element is neither reliable nor consistent. Errors are frequent—cells continually die—and reliability is maintained through massive redundancy as well as through the inherently error-tolerant nature of the computational process and, for that matter, the relatively high error-tolerance of the resulting behavior.

These last points cannot be overemphasized. The body, the brain, and human social interaction have all coevolved to tolerate large variations in performance under a wide-ranging set of environmental conditions. It is a remarkably error-tolerant and forgiving system. It uses both electrical and chemical systems of communication and processing. Conscious and subconscious processing probably use different computational mechanisms, and the role of emotions and affect is not yet understood.

Human language serves as a good example of the evolution of a robust, redundant, and relatively noise-insensitive means of social communication. Errors are corrected so effortlessly that often neither party is aware of the error or the correction. Communication relies heavily upon a shared knowledge base, intentions, and goals; people with different cultural backgrounds often clash, even though they speak the same language. The result is a marvelously complex structure for social interaction and communication. Children learn language without conscious effort, yet the complexities of human language still defy complete scientific understanding.

Biological versus Technological Evolution

We humans have evolved to fit the natural environment. At the same time we have learned to modify and change the environment. This process, in which we've changed to fit the world while simultaneously changing the world, leads to further evolutionary change. Until recently, this coevolution proceeded at a human pace. We developed language and tools. We discovered how to control fire and construct simple tools. The tools became more complex as simple tools became machines. The process was slow, the better to fit the new ways with the old, the new methods with human capabilities.

Biological evolution of humankind proceeds too slowly to be visible, but there is a kind of technological and environmental evolution that proceeds rapidly. We evolve our human-made artifacts to fit our abilities. This evolution is similar to, yet different from, the biological kind. For one thing, it has a history: It is Lamarckian, in that lessons learned in one generation can be propagated to future ones. Nonetheless, it is an evolutionary process, because it tends to be unguided except by rules of survival. Each new generation is but a small modification of the previous one.

A good illustration of how an evolutionary process shapes our human-invented artifacts is sports. Sports require an exquisite mix of the doable and the difficult. Make a game too easy and it loses its appeal. Make it too difficult and it is unplayable. The range from too easy to too difficult is huge, and fortunately so. One of our traits is the ability to learn, to develop skills far beyond that which the unpracticed person can do. Thus, some games, such as tic-tac-toe, which seem difficult when first encountered, are so readily mastered that they soon become boring. A successful game is one that has a wide range of complexity, playable by beginners and experts alike, although not necessarily at the same time. Successful games include soccer, rugby, tennis, basketball, baseball, football, chess, go, checkers, poker, and bridge. These are multidimensional, rich, and multifaceted. As a result, the beginner can enjoy part of their charm while the expert can exploit all the multiple dimensions.

Games work well when they do not use too much technology. The reason is simple: Games are suited to human reaction times, size, and strength. Add too much technology to the mix, and you soon move the game beyond the reach of human abilities. This is aptly illustrated in war, the deadly dueling exercises in which the armies of the world pit themselves one against the other. But here, the technologies are deliberately exploited to exceed human capability, so much so that it can take ten years of training to master a modern jet fighter plane, and even then the human pilot is rendered temporarily unconscious during violent maneuvers. These are games not fit for people.

Alas, the slow, graceful coevolution of people and environment, and of the tools, artifacts, and games that we have designed, no longer holds. Each generation benefits from the one before, and the accumulated knowledge leads to more rapid change. We benefit greatly from this cumulative buildup of knowledge, but the price we pay is that each succeeding generation has more and more to learn. The result is that the past acts both as a wonderful starting point, propelling us forward on the shoulders of giants, and as a massive anchor, compelling us to spend more and more time at school, learning the accumulated wisdom of the ages, to the point that one's motivation and energy may be depleted before the studies are over.

The Ever-Increasing Pace of Change

Once upon a time it was possible for people to learn a great deal about their culture. After all, things changed slowly, at a human pace. As they grew up, children learned about what had happened before, and from then on, they could keep up with the rate of change. The technology changed slowly. Moreover, it was mechanical, which meant it was visible. Children could explore it. Teenagers could disassemble it. Young adults could hope to improve it.

Once upon a time technological evolution proceeded at a human pace. Crafts and sports evolved over a lifetime. Even though the results could be complex, the reason behind the complexity could usually be seen, examined, and talked about. The technology could be lived with and experienced. As a result, it could be learned.

Today, this is no longer possible. The slow evolutionary pace of life is no longer up to the scale and pace of technological change. The accumulation of knowledge is enormous, increasing with every passing year. Once upon a time, a few years of schooling—or even informal learning—was sufficient. Today, formal schooling is required, and the demands upon it continually increase. The number of different topics that must be mastered, from history and language to science and technology to practical knowledge and skills, is ever-increasing. Once a grade-school education would suffice for most people. Then high school was required. Then college, postgraduate education, and even further education after that. Today, no amount of education is sufficient.

Scientists no longer are able to keep up with advances even within their own field, let alone in all of science. As a result, we are in the age of specialization, where it is all one person can do to keep up with the pace in some restricted domain of endeavor. But with nothing but specialists, how can we bridge the gaps?

The new technologies can no longer be learned on their own. Today, technology tends to be electronic, which means that its operation is invisible, for it takes place inside of semiconductor circuits through the transfer of voltages, currents and electromagnetic fields, all of which are invisible to the eye. A single computer chip may have ten million components, and chips with 100 million components are in the planning stage. Who could learn such things by disassembly, even were disassembly possible? So, too, with computer programs; a program with hundreds of thousands of lines of instructions is commonplace, and some have millions of lines.

Worse, the new technology can often be arbitrary, inconsistent, overly complex, and irrelevant. It is all up to the whim of the designer. In the past, physical structures posed their own natural constraints upon the design and the resulting complexity. But with information technologies, the result can be as simple or complex as the designer wills it to be, and far too few designers have sufficient appreciation for the requirements of the people who must use their designs.

Even when a designer is considerate of the users of the technology, there may be no natural relationship between one set of designs and another. In the physical world, the natural constraints of physical objects meant that similar tools worked in similar ways. Not so in the world of information: Very similar tools may work in completely different—perhaps even contradictory—ways.

Treating People like Machines

What an exciting time the turn of the century must have been! The period from the late 1800s through the early 1900s was one of rapid change, in many ways paralleling the changes that are taking place now. In a relatively short period of time, the world went through rapid, almost miraculous technological invention, forever changing the lives of its citizens, society, business, and government. In this period, the incandescent light was developed and electric power plants sprung up across the nation. Electric motors were developed to power factories. The telegraph spanned the American continent and the world, followed by the telephone. With the phonograph, for the first time in history voices, songs, and sounds could be preserved and replayed at will. At the same time, mechanical devices were increasing in power. The railroad was rapidly expanding its coverage. Steam-powered ships traveled the oceans. The automobile was invented, first as expensive, hand-made machines, starting with Daimler and Benz in Europe. Henry Ford developed the first assembly line for the mass-production of relatively

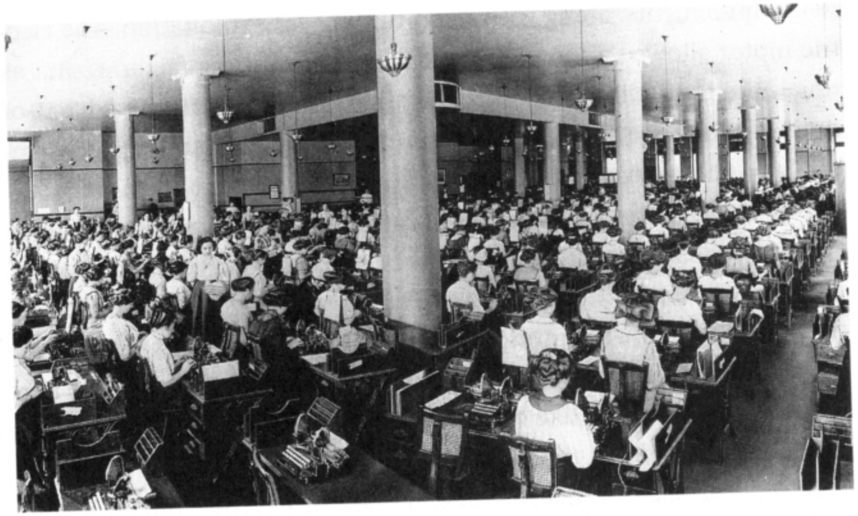

Figure 7.2
Typical pre–World War I office in a large American mail-order house, c. 1912.
(Photograph courtesy of Corbis-Bettmann.)

inexpensive automobiles. The first airplane was flown and within a few decades would carry mail, passengers, and bombs. Photography was common and motion pictures were on the way. Radio was soon to come, allowing signals to be sent all across the world without the need for wire. It was a remarkable period of change.

It is difficult today to imagine life without these products of technology. At night the only lighting was through flames: candles, fireplaces, oil and kerosene lamps, and in some places, gas. Letters were the primary means of communication, and although mail delivery within a large city was rapid and efficient, with delivery offered more than once each day, delivery across distances could take days or even weeks. Travel was difficult, and many people never ventured more than 30 miles from their homes during their entire lives. But in what to a historian is a relatively short period, the world changed dramatically in ways that affected everyone—not just the rich and upper class, but all levels of society.

Light, travel, entertainment: All changed through human inventions. Work did too, although not always in beneficial ways. The factory already existed, but the new technologies and processes brought forth new requirements, along with opportunities for exploitation. The electric motor allowed a more efficient means of running factories. But as usual, the largest impact was social and organizational: the advent of time-and-motion studies, of "scientific management," and of the assembly line. These developments analyzed human work patterns into a series of small actions. The belief was that if each action could be standardized, each organized into "the one best way,"[3] then automated factories could reap the benefits of even greater efficiency and productivity. The consequence was dehumanization of the worker. Now the worker was considered to be just another machine in the factory, analyzed like one, treated like one, and asked not to think on the job, for thinking slowed down the action.

The era of mass production and the assembly line resulted in part from the efficiencies of the "disassembly line" developed by the meat-packing factories. The tools of scientific management took into account the physical properties of the human body but overlooked the mental and psychological ones. The result was to cram ever more motions into the working day, treating the factory worker as a cog in a machine, deliberately depriving work of its meaning, all in the name of efficiency. These beliefs have stuck with us, and although today we do not go to quite the extremes advocated by the early practitioners of scientific management, the die was cast for the mindset of ever-increasing efficiency, ever-increasing productivity from the workforce. The principle of improved efficiency is hard to disagree with. The question is, at what price?[4]

The work of Frederick Taylor, some people believe, has had a larger impact upon the lives of people in this century than that of anyone else. He thought there was "the one best way" of doing things. His book, *The Principles of Scientific Management*,[5] published in 1911, guided factory development and workforce habits throughout the world, from those in the United States to Stalin's attempt to devise an efficient communist workplace in the newly formed Soviet Union. He is primarily responsible for our notions of efficiency and of the work practices followed in industry around the world, and even for the sense of guilt we sometimes feel when we have been "goofing off" instead of attending to business.

Taylor's "scientific management" was a detailed, careful study of jobs, breaking down each task into its basic components. Once you knew the components, you could devise the most efficient way of doing things, devise procedures that enhanced performance and increased the efficiency of workers. If Taylor's methods were followed properly, management could raise workers' pay while increasing company profit. In fact, Taylor's methods required management to raise the pay, for money was used as the incentive to get workers to follow the procedures. According to Taylor, everybody would win: The workers would get more money, the management more production and more profit. Sounds wonderful, doesn't it? The only problem was that workers hated it.

Taylor, you see, thought of people as simple machines. Once management found the best way to do things, it should have its workers do it that way, hour after hour, day after day. Efficiency permitted no deviation. Thought was eliminated. According to Taylor, the sort of people who could shovel dirt, do simple cutting, lathing, and drilling, those who perform the lowest-level tasks, were not capable of thought. He regarded them as "brute laborers." Furthermore, if thought was called for, there must be some lack of clarity in the procedures or the process, which signaled that the procedures were wrong. The problem with thinking, explained Taylor, was not only that most workers were incapable of it, but that it slowed the work down. That's certainly true: Why, if we never had to think, just imagine how much faster we could work. In order to eliminate the need for thought, Taylor stated that it was necessary to reduce all work to the routine:[6] That is, all work except for people like him who didn't have to keep fixed hours, who didn't have to follow procedures, who were paid hundreds of times greater wages than the so-called brutes, and who were allowed, even encouraged, to think.

Taylor thought that the world was neat and tidy. If only everyone would do things according to procedure, everything would run smoothly, producing a clean, harmonious world. Taylor may have thought he understood machines, but he certainly didn't understand people. In fact, he didn't really understand the complexity of machines or of work. And he certainly didn't understand the complexity of the world.

The World Is Not Neat and Tidy

Not only don't things always work as planned, but the notion of "plan" itself is suspect. Organizations spend a lot of time planning, but although the act of planning is useful, the plans themselves are often obsolete even before their final printing.

There are many reasons for this. Those philosophically inclined can talk about the fundamental nature of quantum uncertainty, of the fundamental statistical nature of matter. Alternatively, one can talk of complexity theory and chaos theory, where tiny perturbations can have major, unexpected results. I prefer to think of the difficulties as consequences of the complex interactions that take place among the trillions

of events and objects in the world, so many interactions that even if science were advanced enough to understand each individual one there are simply too many combinations and variations possible ever to have worked out all possibilities. All of these views are compatible with one another.

Consider these examples of complex situations in which things habitually go wrong:

- A repair crew disconnects a pump from service in a nuclear power plant,[7] carefully placing tags on the controls so that the operators will know the particular unit is temporarily out of service. Later, as the operators attempt to deal with an unrelated problem, they initially diagnose it in an erroneous, albeit reasonable, way. Eventually, the problem becomes so serious that the entire plant is destroyed in the worst accident in the history of American nuclear power. Among the factors hindering the correct diagnosis of the situation is that the tags so carefully placed to indicate the out-of-service unit hang in front of another set of indicators, blocking them from view. Could this have been predicted beforehand? Maybe. But it wasn't.
- A hospital x-ray technician enters a dosage for an x-ray machine, then corrects the setting after realizing the machine is in the wrong mode.[8] The machine's computer program, however, wasn't designed to handle a rapidly made correction, so it did not properly register the new value. Instead, it delivered a massive overdose to the patient. Sometime later, the patient died as a result. The accident goes undiagnosed, because as far as anyone can determine, the machine had performed correctly. Moreover, the effect of overdosage doesn't show up immediately, so when the symptoms were reported, they were not correlated with the incident, or for that matter, with the machine. When the machine's performance first comes under suspicion, the company who manufactured it explains in detail why such an accident is impossible. The situation repeats itself in several different hospitals, killing a number of patients before a sufficient pattern emerges, so that the problem can be recognized and the design of the machine fixed. Could this have been predicted beforehand? Maybe. But it wasn't.

- The French air-traffic controllers seem to be forever complaining, frequently calling strikes and protests. American air traffic controllers aren't all that happy either. And guess what the most effective protest method is? Insisting on following procedures. On normal days, if the workers follow the procedures precisely, work slows up, and in the case of air-traffic control, airline traffic around the entire world is affected. The procedures must be violated to allow the traffic to flow smoothly. Of course, if there is an accident and the workers are found not to have followed procedures, they are blamed and punished.
- The United States Navy has a formal, rigid hierarchy of command and control, with two classes of workers—enlisted crew and officers—and a rigid layer of formal rank and assignment. There are extensive procedures for all tasks. Yet in their work habits, especially in critical operations, rank seems to be ignored and crew members frequently question the actions. Sometimes they even debate the appropriate action to be taken. The crew, moreover, is always changing. There are always new people who have not learned the ship's procedures, and even the veterans often don't have more than two or three year's experience with the ship; the Navy has a policy of rotating assignment. Sounds horrible, doesn't it? Isn't the military supposed to be the model of order and structure? But wait. Look at the outcome: The crew functions safely and expertly in dangerous, high-stress conditions. What is happening here?

These examples illustrate several points. The world is extremely complex, too complex to keep track of, let alone predict. In retrospect, looking back after an accident, the problem seems obvious. There are usually a few simple actions that, had they been taken, would have prevented the accident, precursor events that, had they been perceived and interpreted properly, would have provided sufficient warning. Sure, but this is hindsight after we know how things turned out.

Remember, life is complex. Lots of stuff is always happening, most of which is irrelevant to the task at hand. We all know that it is important to ignore the irrelevant and attend to the relevant. But how does one know which is which?

We humans are a complex mixture of motives and mechanisms. We are sense-making creatures, always trying to understand and give explanations for the things we encounter. We are social animals, seeking company, working well in small groups. Sometimes this is for emotional support, sometimes for assistance, sometimes for selfish reasons, so we have someone to feel superior to, to show off to, to tell our problems to. We are narcissistic and hedonistic, but also altruistic. We are lots of things, sometimes competing, conflicting things. And we are also animals, with complex biological drives that strongly affect behavior: emotional drives, sexual drives, hunger drives. Strong fears, strong desires, strong phobias, and strong attractions.

Making Sense of the World

If an airplane crashes on the border between the United States and Canada, killing half the passengers, in which country should the survivors be buried?

People try to make sense of the world. We assume that information is sensible, and we do the best we can with what we receive. This is a virtue. It makes us successful communicators, efficient and robust in going about our daily activities. It also means we can readily be tricked. It wasn't Moses who brought the animals aboard the ark, it was Noah. It isn't the survivors who should be buried, it is the casualties.

It's a good thing we are built this way: This compliance saves us whenever the world goes awry. By making sense of the environment, by making sense of the events we encounter, we know what to attend to, what to ignore. Human attention is the limiting factor, a well-known truism of psychology and of critical importance today. Human sensory systems are bombarded with far more information than can be processed in depth; some selection has to be made. Just how this is done has been the target of prolonged investigation by numerous cognitive scientists who have studied people's behavior when overloaded with information, by neuroscientists who have tried to follow the biological processing of sensory signals, and by a host of other investigators. I was one of them: I spent almost ten years of my research career studying the mechanisms of human attention.

One useful tool to deepen our understanding of the cognitive process of attention comes from the notion of a "conceptual model." This idea has great importance in chapter 8 when I discuss how to design technology that people can use. A conceptual model is, to put it simply, a story that makes sense of a situation.

I sit at my desk with a large number of sounds impinging upon me. It is an easy matter to classify the sounds. What is all that noise outside? A family must be riding bicycles and the parents are yelling to their children. And the neighbor's dogs are barking at them, which is why my dogs started barking. Do I really know this? No. I didn't even bother to look out the window. My mind subconsciously, automatically created the story, creating a comprehensive explanation for the noises, even as I concentrated upon the computer screen.

How do I know what really happened? I don't. I listened to the sounds and created an explanation, one that was logical, heavily dependent upon past experience with those sound patterns. It is most likely correct, but I don't really know.

Note that the explanation also told me which sounds went together. I associated the barking dogs with the family of bicyclists. Maybe the dogs were barking at something else. Maybe. How do I know they were on bicycles? I don't, but it is a common activity near my home, which means it is reasonably likely. The point is not that I might be wrong, the point is that this is normal human behavior. Moreover, it is human behavior that stands us in good stead. I am quite confident that my original interpretations were correct, confident enough that I won't bother to check. I could be wrong.

A good conceptual model of events allows us to classify them into ones that are relevant, and ones that are not, dramatically simplifying life. We attend to the relevant and only monitor the irrelevant. Mind you, this monitoring and classification are completely subconscious. The conscious mind is usually unaware of the process. Indeed, the whole point is to reserve the conscious mind for the critical events of the task being attended to by suppressing most of the other irrelevant events.

On the whole, human consciousness avoids paying attention to the routine. Conscious processing attends to the nonroutine, to discrepancies and novelties, to things that go wrong. As a result, we are sensitive to changes in the environment, remarkably insensitive to the commonplace, the routine.

Most of the time people do brilliantly. People are very good at predicting things. Experts are particularly good at this because of their rich experience. When a particular set of events occurs, they know exactly what will follow.

But what happens when the unexpected happens? Do we go blindly down the path of the most likely interpretation? In fact, this is the recommended strategy and will usually lead to a correct diagnosis. Most of the time we can then find a solution. You seldom hear about those instances. Headlines appear when things go wrong, not when they go right.

Look back at the incidents I described earlier. Consider the role played by conceptual models in the nuclear power incident, the famous Three Mile Island event that destroyed the power-generating unit and caused such a public loss in confidence in nuclear power that no American plant has been built since. The operators' conceptual model of the events led them to misdiagnose the situation, leading to a major calamity. But the misdiagnosis was a perfectly reasonable one. As a result, they concentrated on items they thought relevant to their diagnosis and missed other cues, which they thought were just part of the normal background noise. The tags that blocked the view would not normally have been important.

In the hospital x-ray situation, the real error was in the design of the software system, but even here, the programmer erred in not thinking through all of the myriad possible sequences of operation, something not easy to do. There are better ways of developing software[9] that would have made it more likely to have caught these problems before the system was released to hospitals, but even then, there are no guarantees. As for the hospital personnel who failed to understand the relationship, well, they too were doing the best they could to interpret the events and to get through their crowded, hectic days. They interpreted things according to normal events, which was wrong only because this one was abnormal.

Do we punish people for failure to follow procedures? This is what Frederick Taylor would have recommended. After all, management determines the one best way to do things, writes a detailed procedure to be followed in every situation, and expects workers to follow them. That's how we get maximum efficiency. But how is it possible to write a procedure for absolutely every possible situation, especially in a world filled with unexpected events? Answer: It's impossible.

Procedures and rule books dominate industry. The rule books take up huge amounts of shelf space. In some industries, it is impossible for any individual to know all the rules. The situation is made even worse by legislatures that can't resist adding new rules. Was there a major calamity? Pass a law prohibiting some behavior, or requiring some other behavior. Of course, the law strikes at the factor easiest to blame, whereas in most complex situations, multiple factors interact and no single one is fully responsible. Nonetheless, new rules are written, controlling sense and reasonableness in the conduct of business.

Do we need procedures? Of course. But procedures must be designed with care and attention to the social, human side of the operation. The best procedures will mandate outcomes, not methods. Methods change; it is the outcomes we care about. Remember the striking air traffic controllers who brought things to a halt following procedures? The same condition exists in most industries. If the procedures are followed exactly, work slows to an unacceptable level. In order to perform properly it is necessary to violate the procedures. Workers get fired for lack of efficiency, which means they are subtly, unofficially encouraged to violate the procedures. Unless something goes wrong, in which case they can be fired for failure to follow the procedures.

Now look at the Navy. The constant critiques and arguments are not what they seem to be. The apparent chaos is a carefully honed system, tested and evolved over generations, that maximizes safety and efficiency in the face of numerous mistakes, novel circumstances, and a wide range of skills and knowledge among the crew. Having everyone participate and question the actions serves several roles simultaneously. The very ambiguity, the continual questioning and debate keeps everyone in touch with the activity, thereby providing redundant checks on the actions. This adds to the safety, for now it is likely for errors to get detected before they have caused problems. The newer crew members have a lot to learn, and the public discussions among the other crew serve as valuable training exercises, not in some abstract fashion, but in situations where it really matters. And by not punishing people when they speak out, question, or even bring the operations to a halt, they encourage continual learning and performance enhancement. It makes for an effective, well-tuned team.

New crew members don't have the experience of older ones. This means they are not efficient, don't always know what to do, and perform slowly. They need a lot of guidance. The system automatically provides this constant supervision and coaching, allowing people to learn on the job. At the same time, because the minds of the new crew members are not yet locked into the routines, their questioning can sometimes reveal errors. Their fresh approach challenges the official mindset, asking whether the generally accepted explanation of events is correct. This is the best way to avoid errors of misdiagnosis.

The continual challenge to authority goes against conventional wisdom and is certainly a violation of the traditional hierarchical management style. But it is so important to safety that the aviation industry now has special training in crew management, where the junior officers in the cockpit are encouraged to question the actions of the captain. In turn, the captain, who used to be thought of as the person in command, with full authority and never to be questioned, is now trained to encourage crew members to question every act. The end result may look less disciplined, but it is far safer.

The Navy's way of working is sensible. Accidents are minimized. Despite the fact that the Navy is undertaking dangerous operations under periods of fast pace and high stress, there are remarkably few mishaps. If the Navy would follow formal procedures and a strict hierarchy of rank, the result would very likely be an increase in the accident rate.[10] Other industries would do well to copy this behavior. Fred Taylor would turn over in his grave (efficiently, without any wasted motion).

Human Error

Machines, including computers, don't err, in the sense that they are fully deterministic, always returning the same value for the same inputs and operations. Someday we may have stochastic or quantum computation, but even then we will expect them to follow precise laws of operation. When computers do err, it is either because a part has failed or because of human error, either in design specification, programming, or faulty construction. People are not fully deterministic; ask a person to repeat an operation, and the repetition is subject to numerous variations.

People do err, but primarily because they are asked to perform unnatural acts: to do detailed arithmetic calculations, to remember details of some lengthy sequence or statement, or to perform precise repetitions of actions, all the result of the artificial nature of invented artifacts. They err when attempting to alter habitual behavior, forgetting to mail a letter on the way to work or stop at the store on the way home. Slips of the tongue are common, although often the intended meaning is still conveyed sufficiently well that the errors are not even noticed by either speaker or listener. People leave jackets on airplanes and babies on buses. They lock themselves out of home and car. People are expert at making errors.

Human error matters primarily because we followed a technology-centered approach in which it matters. A human-centered approach would make the technology robust, compliant, and flexible. The technology should conform to the people, not people to the technology.

Human languages provide an excellent example of how systems can be tailored for human capabilities, providing a rich structure for communication and social interaction while being extremely tolerant of error. Language is so natural to learn that it is done without any formal instruction: only severe brain impairment can eliminate the capability of learning language. Note that "natural" does not mean "easy"; it takes ten to fifteen years to master one's native language. Second language learning can be excruciatingly difficult.

Natural language, unlike programming language, is flexible, ambiguous, and heavily dependent on shared understanding, a shared knowledge base, and shared cultural experiences. Errors in speech are seldom important: Utterances can be interrupted, restarted, even contradicted, with little difficulty in understanding. The system makes natural language communication extremely robust.

Today, when faced with human error, the traditional response is to blame the human and institute a new training procedure: blame and train. But when the vast majority of industrial accidents is attributed to human error, it indicates that something is wrong with the system, not the people. Consider how we would approach a system failure caused by a noisy environment: We wouldn't blame the noise, we would instead design a system that was robust in the face of noise.

This is exactly the approach that should be taken in response to human error: redesign the system to fit the people who must use it. This means avoiding the incompatibilities between human and machine that generate error, making it so that errors can be rapidly detected and corrected, and being tolerant of error. To blame and train does not solve the problem.

Humans and Computers as Complementary Systems

Because humans and computers are such different kinds of systems, it should be possible to develop a strategy for complementary interaction. Alas, today's approaches are wrong. One major theme is to make computers more like humans. This is the original dream behind classic artificial intelligence: to simulate human intelligence. Another theme is to make people more like computers. This is how technology is designed today; the designers determine the needs of the technology and then ask people to conform to those needs. The result is an ever-increasing difficulty in learning the technology, and an ever-increasing error rate. It is no wonder that society exhibits an ever-increasing frustration with technology.

Consider the following attributes of humans and machines presented from today's machine-centered point of view:[11]

The Machine-Centered View

People	Machines
Vague	Precise
Disorganized	Orderly
Distractible	Undistractible
Emotional	Unemotional
Illogical	Logical

Note how the humans lose: All the attributes associated to people are negative, all the ones associated with machines are positive. But now consider attributes of humans and machines presented from a human-centered point of view:

The Human-Centered View

People	Machines
Creative	Unoriginal
Compliant	Rigid
Attentive to change	Insensitive to change
Resourceful	Unimaginative

Now note how machines lose: all the attributes associated with people are positive, all the ones associated with machines are negative.

The basic point is that the two different viewpoints are complementary. People excel at qualitative considerations, machines at quantitative ones. As a result, for people, decisions are flexible because they follow qualitative as well as quantitative assessment, modified by special circumstances and context. For the machine, decisions are consistent, based upon quantitative evaluation of numerically specified, context-free variables. Which is to be preferred? Neither: We need both.

It's good that computers don't work like the brain. The reason I like my electronic calculator is that it is accurate; it doesn't make errors. If it

were like my brain, it wouldn't always get the right answer. This very difference is what makes the device so valuable. I think about the problems and the method of attack. It does the dull, dreary details of arithmetic—or in more advanced machines, of algebraic manipulation and integration. Together, we are a more powerful team than either of us alone.

The same principle applies to all our machines; we should capitalize on the difference, for together we complement one another. This is useful, however, only if the machine adapts itself to human requirements. Alas, most of today's machines, especially the computer, force people to use them on their terms, terms that are antithetical to the way people work and think. The result is frustration, an increase in the rate of error (usually blamed on the user—human error—instead of on faulty design), and a general turning away from technology.

Will the interactions between people and machines be done correctly in the future? Might schools of computer science start teaching the human-centered approach that is necessary to reverse the trend? I don't see why not.