CS247 Reading 1 - Error

From:
Don Norman
The Design of Everyday Things
Basic Books Paperback Edition
ISBN: 0465067107

TO ERR IS HUMAN

"LONDON—An inexperienced computer-operator pressed the wrong key on a terminal in early December, causing chaos at the London Stock Exchange. The error at stockbrokers Greenwell Montagu led to systems staff working through the night in an attempt to cure the problem."[1]

People make errors routinely. Hardly a minute of a normal conversation can go by without a stumble, a repetition, a phrase stopped midway through to be discarded or redone. Human language provides special mechanisms that make corrections so automatic that the participants hardly take notice; indeed, they may be surprised when errors are pointed out. Artificial devices do not have the same tolerance. Push the wrong button, and chaos may result.

Errors come in several forms. Two fundamental categories are slips and mistakes. Slips result from automatic behavior, when subconscious actions that are intended to satisfy our goals get waylaid en route. Mistakes result from conscious deliberations. The same processes that make us creative and insightful by allowing us to see relationships between apparently unrelated things, that let us leap to correct conclusions on the basis of partial or even faulty evidence, also lead to error.

Our ability to generalize from small amounts of information helps tremendously in new situations; but sometimes we generalize too rapidly, classifying a new situation as similar to an old one when, in fact, there are significant discrepancies. False generalizations can be hard to discover, let alone eliminate.

The differences between slips and mistakes are readily apparent in the analysis of the seven stages of action. Form an appropriate goal but mess up in the performance, and you've made a slip. Slips are almost always small things: a misplaced action, the wrong thing moved, a desired action undone. Moreover, they are relatively easy to discover by simple observation and monitoring. Form the wrong goal, and you've made a mistake. Mistakes can be major events, and they are difficult or even impossible to detect—after all, the action performed is appropriate for the goal.

Slips

A colleague reported that he went to his car to drive to work. As he drove away, he realized that he had forgotten his briefcase, so he turned around and went back. He stopped the car, turned off the engine, and unbuckled his wristwatch. Yes, wristwatch, instead of his seatbelt.

Most everyday errors are slips. Intend to do one action, find yourself doing another. Have a person say something clearly and distinctly to you, but "hear" something quite different. The study of slips is the study of the psychology of everyday errors—what Freud called "the psychopathology of everyday life." Some slips may indeed have hidden, darker meanings, but most are accounted for by rather simple events in our mental mechanisms.[2]

Slips show up most frequently in skilled behavior. We don't make so many slips in things we are still learning. In part, slips result from a lack of attention. On the whole, people can consciously attend to only one primary thing at a time. But we often do many things at once. We walk while we talk; we drive cars while we talk, sing, listen to the radio, use a telephone, take notes, or read a map. We can do more than one thing at a time only if most of the actions are done automatically, subconsciously, with little or no need for conscious attention.

Doing several things at once is essential even in carrying out a single task. To play the piano, we must move the fingers properly over the keyboard while reading the music, manipulating the pedals, and listening to the resulting sounds. But to play the piano well, we should do these things automatically. Our conscious attention should be focused on the higher levels of the music, on style, and on phrasing. So it is with every skill. The low-level, physical movements should be controlled subconsciously.

TYPES OF SLIPS

Some slips result from the similarities of actions. Or an event in the world may automatically trigger an action. Sometimes our thoughts and actions may remind us of unintended actions, which we then perform. We can place slips into one of six categories: capture errors, description errors, data-driven errors, associative activation errors, loss-of-activation errors, and mode errors.

CAPTURE ERRORS

"I was using a copying machine, and I was counting the pages. I found myself counting '1, 2, 3, 4, 5, 6, 7, 8, 9, 10, Jack, Queen, King.' I have been playing cards recently."[3]

Consider the common slip called the capture error, in which a frequently done activity suddenly takes charge instead of (captures) the one intended.[4] You are playing a piece of music (without too much attention) and it is similar to another (which you know better); suddenly you are playing the more familiar piece. Or you go off to your bedroom to change your clothes for dinner and find yourself in bed. (This slip was first reported by William James in 1890.) Or you finish typing your thoughts on your word processor or text editing program, turn off the power, and go off to other things, neglecting to save any of your work. Or you get into your car on Sunday to go to the store and find yourself at the office.

The capture error appears whenever two different action sequences have their initial stages in common, with one sequence being unfamiliar and the other being well practiced. Seldom, if ever, does the unfamiliar sequence capture the familiar one.

DESCRIPTION ERRORS

A former student reported that one day he came home from jogging, took off his sweaty shirt, and rolled it up in a ball, intending to throw *it in the laundry basket. Instead he threw it in the toilet. (It wasn't poor aim: the laundry basket and toilet were in different rooms.)*

In the common slip known as the description error, the intended action has much in common with others that are possible. As a result, unless the action sequence is completely and precisely specified, the intended action might fit several possibilities. Suppose that my tired student in the example formed a mental description of his intended action something like "throw the shirt into the opening at the top of the container." This description would be perfectly unambiguous and sufficient were the laundry basket the only open container in sight; but when the open toilet was visible, its characteristics matched the description and triggered the inappropriate action. This is a description error because the internal description of the intention was not sufficiently precise. Description errors usually result in performing the correct action on the wrong object. Obviously, the more the wrong and right objects have in common, the more likely the errors are to occur. Description errors, like all slips, are more likely when we are distracted, bored, involved in other activities, under extra stress, or otherwise not inclined to pay full attention to the task at hand.

Description errors occur most frequently when the wrong and right objects are physically near each other. People have reported a number of description errors to me.

Two clerks in a department store were both on the telephone to verify credit cards while simultaneously dealing with a customer and filling out a credit card form. One sales clerk had passed in back of the other to reach the charge forms. When this clerk finished preparing the sales slip, she hung up the handset on the wrong telephone, thereby terminating the other clerk's call.

A person intended to put the lid on a sugar bowl, but instead put it on a coffee cup (with the same size opening).

I had a report of someone who planned to to pour orange juice into a glass but instead poured it into a coffee cup (adjacent to the glass).

Another person told me of intending to pour rice from a storage jar into a measuring cup, but instead pouring cooking oil into the measuring cup (both the oil and the rice were kept in glass containers on the counter).

Some things seem designed to cause slips. Long rows of identical switches are perfect setups for description errors. Intend to flip one switch, instead flip a similar-looking one. It happens in industrial plants, aircraft, homes, anywhere. When different actions have similar descriptions, there is a good chance of mishap, especially when the operator is experienced and well practiced and therefore not paying full attention, and if there are more important things to do.

DATA-DRIVEN ERRORS

"I was assigning a visitor a room to use. I decided to call the department secretary to tell her the room number. I used the telephone in the alcove outside the room, with the room number in sight. Instead of dialing the secretary's phone number—which I use frequently and know very well—I dialed the room number."

Much human behavior is automatic, for example, brushing away an insect. Automatic actions are data driven—triggered by the arrival of the sensory data. But sometimes data-driven activities can intrude into an ongoing action sequence, causing behavior that was not intended.

ASSOCIATIVE ACTIVATION ERRORS

"My office phone rang. I picked up the receiver and bellowed 'Come in' at it."[5]

If external data can sometimes trigger actions, so, too, can internal thoughts and associations. The ringing of the telephone and knocking on the door both signal the need to greet someone. Other errors occur from associations among thoughts and ideas. Associative activation errors are the slips studied by Freud; you think something that ought not to be said and then, to your embarrassment, you say it.

LOSS-OF-ACTIVATION ERRORS

"I have to go to the bedroom before I start working in the dining room. I start going there and realize as I am walking that I have no idea why I am going there. Knowing myself, I keep going, hoping that something in the bedroom will remind me. . . . I get there but still cannot recall what I wanted . . . so I go back to the dining room. There I realize that my glasses are dirty. With great relief I go back to the bedroom, get my handkerchief, and wipe my glasses clean."

One of the more common slips is simply forgetting to do something. More interesting is forgetting part of the act, remembering the rest, as in the story above where the goal was forgotten, but the rest of the action continued unimpaired. One of my informants walked all the way through the house to the kitchen and opened the refrigerator door; then he wondered why he was there. Lack-of-activation errors occur because the presumed mechanism—the "activation" of the goals—has decayed. The less technical but more common term would be "forgetting."

MODE ERRORS

"I had just completed a long run from my university to my home in what I was convinced would be record time. It was dark when I got home, so I could not read the time on my stopwatch. As I walked up and down the street in front of my home, cooling off, I got more and more anxious to see how fast I had run. I then remembered that my watch had a built-in light, operated by the upper right-hand button. Elated, I depressed the button to illuminate the reading, only to read a time of zero seconds. I had forgotten that in stopwatch mode, the same button [that in the normal, time-reading mode would have turned on a light] cleared the time and reset the stopwatch."

Mode errors occur when devices have different modes of operation, and the action appropriate for one mode has different meanings in other modes. Mode errors are inevitable any time equipment is designed to have more possible actions than it has controls or displays, so the controls must do double duty. Mode errors are especially likely where the equipment does not make the mode visible, so the user is expected to remember what mode has been established, sometimes for many hours.

Mode errors are common with digital watches and computer systems (especially text editors). Several accidents in commercial aviation can be attributed to mode errors, especially in the use of the automatic pilots (which have a large number of complex modes).

DETECTING SLIPS

Although slips are relatively easy to detect because there is a clear discrepancy between goal and result, detection can only take place if

there is feedback. If the result of the action is not visible, how can a misaction be detected? Even when a mismatch is noted, the person may not believe that the error occurred. Some trail of the sequence of actions that was performed is valuable.

Even when an error has been detected, it may not be clear what the error was.

"Alice" was driving a van and noticed that the rearview mirror on the passenger side was not adjusted properly. Alice meant to say to the passenger on the right, "Please adjust the mirror," but instead said "Please adjust the window."

The passenger, "Sally," was confused and asked, "What should I do? What do you want?"

Alice repeated the request: "Adjust the window for me."

The situation continued through several frustrating cycles of conversation and attempts by the passenger to understand just what adjustments should be made to the window. The error-correction mechanism adopted by the driver was to repeat the erroneous sentence more and more loudly.

In this example, it was easy to detect that something was wrong but hard to discover what. Alice believed the problem was that she couldn't be understood or heard. She was monitoring the wrong part of the action sequence—she had a problem of level.

Actions can be specified at many different levels. Suppose I were driving my car to the bank. At any given moment, the action being performed could be described at many different levels:

- Driving to the bank
- Turning into the parking lot
- Making a right turn
- Rotating the steering wheel clockwise
- Moving my left hand upward and to the right and my right hand downward
- Increasing the tension on the sternocostal portion of the pectoralis major muscle

All these levels are active at the same time. The most global description (the one at the top of the list), is called the high-level specification. The more detailed descriptions, the ones at the bottom of the list, are called the low-level specifications. Any one of them might be in error.

It is often possible to detect that the result of an action is not as planned, but then not to know at which level of specification the error has taken place.

Problems of level commonly thwart the correction of error. My collection of slips includes several examples in which a person detects a problem but attempts to correct it at the wrong level.

One frequent example is the nonworking key, reported to me both for cars and homes. Someone goes to his or her car and the key won't work. The first response is to try again, perhaps holding the key more level or straight. Then the key is reversed, tried upside down. When that fails, the key is examined and perhaps another tried in its stead. Then the door is wiggled, shaken, hit. Finally, the person decides that the lock has broken, and walks around the car to try the other door, at which point it is suddenly clear that this is the wrong car.

In all the situations I have examined the error correction mechanism seems to start at the lowest possible level and slowly works its way higher. Whether this is universally true I do not know, but the hypothesis warrants further examination.

DESIGN LESSONS FROM THE STUDY OF SLIPS

Two different kinds of design lessons can be drawn, one for preventing slips before they occur and one for detecting and correcting them when they do occur. In general, the solutions follow directly from the preceding analyses. For example, mode errors are minimized by minimizing modes, or at least by making modes visible.

Cars provide a number of examples of how design relates to error. A variety of fluids are required in the engine compartment of an automobile: engine oil, transmission oil, brake fluid, windshield washer solution, radiator coolant, battery water. Putting the wrong fluid into a reservoir could lead to serious damage or even an accident. Automobile manufacturers try to minimize these errors (a combination of description and mode errors) by making the different compartments look different—using different shapes and different-size openings—and by adding color to the fluids so that they can be distinguished. Here design by and large prevents errors. But, unfortunately, designers seem to prefer to encourage them.

I was in a taxi in Austin, Texas, admiring the large number of new devices in front of the driver. No more simple radio. In its place was a computer display, so that messages from the dispatcher were now printed on the screen. The driver took great delight in demonstrating all the features to me. On the radio transmitter I saw four identical-looking buttons laid out in a row.

"Oh," I said, "you have four different radio channels."

"Nope," he replied, "three. The fourth button resets all the settings. Then it takes me thirty minutes to get everything all set up properly again."

"Hmm," I said, "I bet you hit that every now and then by accident."

"I certainly do," he replied (in his own unprintable words).

In computer systems, it is common to prevent errors by requiring confirmation before a command will be executed, especially when the action will destroy a file. But the request is ill timed; it comes just after the person has initiated the action and is still fully content with the choice. The standard interaction goes something like this:

USER: Remove file "My-most-important-work."

COMPUTER: Are you certain you wish to remove the file "My-most-important-work"?

USER: Yes.

COMPUTER: Are you certain?

USER: Yes, of course.

COMPUTER: The file "My-most-important-work" has been removed.

USER: Oops, damn.

The user has requested deletion of the wrong file but the computer's request for confirmation is unlikely to catch the error; the user is confirming the action, not the file name. Thus asking for confirmation cannot catch all slips. It would be more appropriate to eliminate irreversible actions: in this example, the request to remove a file would be handled by the computer's moving the file to some temporary holding place. Then the user would have time for reconsideration and recovery.

At a research laboratory I once directed, we discovered that people would frequently throw away their records and notes, only to discover the next day that they needed them again. We solved the problem by getting seven trash cans and labeling them with the days of the week.

Then the trash can labeled Wednesday would be used only on Wednesdays. At the end of the day it was safely stored away and not emptied until the next Tuesday, just before it was to be used again.

People discovered that they kept neater records and books because they no longer hesitated to throw away things that they thought would probably never be used again; they figured it was safe to throw something away, for they still had a week in which to change their minds.

But design is often a tradeoff. We had to make room for the six reserve wastebaskets, and we had a never-ending struggle with the janitorial staff, who kept trying to empty all of the wastebaskets every evening. The users of the computer center came to depend upon the "soft" nature of the wastebaskets and would discard things that they otherwise might have kept for a while longer. When there was an error—sometimes on the part of the janitorial staff, sometimes on our part in cycling the wastebaskets properly—then it was a calamity. When you build an error-tolerant mechanism, people come to rely upon it, so it had better be reliable.

Mistakes as Errors of Thought

Mistakes result from the choice of inappropriate goals. A person makes a poor decision, misclassifies a situation, or fails to take all the relevant factors into account. Many mistakes arise from the vagaries of human thought, often because people tend to rely upon remembered experiences rather than on more systematic analysis. We make decisions based upon what is in our memory; memory is biased toward overgeneralization and overregularization of the commonplace and overemphasis on the discrepant.

SOME MODELS OF HUMAN THOUGHT

Psychologists have chronicled the failures of thought, the nonrationality of real behavior. Even simple tasks can sometimes throw otherwise clever people into disarray. Even though principles of rationality seem as often violated as followed, we still cling to the notion that human thought should be rational, logical, and orderly. Much of law is based upon the concept of rational thought and behavior. Much of economic theory is based upon the model of the rational human who attempts to optimize personal benefit, utility, or comfort. Many scientists who study artificial intelligence use the mathematics of formal logic—the predicate calculus—as their major tool to simulate thought.

But human thought—and its close relatives, problem solving and planning—seem more rooted in past experience than in logical deduction. Mental life is not neat and orderly. It does not proceed smoothly and gracefully in neat, logical form. Instead, it hops, skips, and jumps its way from idea to idea, tying together things that have no business being put together; forming new creative leaps, new insights and concepts. Human thought is not like logic; it is fundamentally different in kind and in spirit. The difference is neither worse nor better. But it is the difference that leads to creative discovery and to great robustness of behavior.

Thought and memory are closely related, for thought relies heavily upon the experiences of life. Indeed, much problem solving and decision making takes place through attempts to remember some previous experience that can serve as a guide for the present. There have been many theories of human memory. For example, every method of filing things has shown up somewhere along the line as a model for human memory. Do you file photographs neatly in a scrapbook? One theory of memory has postulated that our experiences are neatly encoded and organized, as if in a photo album. This theory is wrong. Human memory is most definitely not like a set of photographs or a tape recording. It mushes things together too much, confuses one event with another, combines different events, and leaves out parts of individual events.

Another theory is based on the filing cabinet model, wherein there are lots of cross references and pointers to other records. This theory has a good deal going for it, and it is probably a reasonable characterization of the most prominent approach today. Of course, it is not called a file cabinet theory. It goes by the names of "schema theory," "frame theory," or sometimes "semantic networks" and "propositional encoding." The individual file folders are defined in the formal structure of the schemas or frames, and the connections and associations among the individual records make the structure into a vast and complex network. The essence of the theory consists of three beliefs, all reasonable and supported by considerable evidence: (1) that there is logic and order to the individual structures (this is what the schema or frame is about); (2) that human memory is associative, with each schema pointing and referring to multiple others to which it is related or that help define the easy to make wrong settings, or to misread an instrument, or to misclassify an event. Design of the social structure that makes false reporting of danger punishable. Turn a nuclear power plant off by mistake and you cost the company hundreds of thousands of dollars; you'll probably lose your job. Fail to turn it off when there is a real incident, and you might lose your life. If you refuse to fly a crowded airliner because the weather looks bad, the company loses lots of money and the passengers get very angry. Take off under those situations and most of the time it works out fine, which encourages risk taking. But every so often there is a disaster.

Tenerife, the Canary Islands, in 1977. A KLM Boeing 747 that was taking off crashed into a Pan American 747 that was taxiing on the runway, killing 583 people. The KLM plane should not have tried to take off then, but the weather was starting to get bad, and the crew had already been delayed for too long (even being on the Canary Islands was a diversion from the scheduled flight—they had to land there because bad weather had prevented them from landing at their scheduled destination); they had not received clearance to take off. And the Pan American flight should not have been on the runway, but there was considerable misunderstanding between the pilots and the air traffic controllers. Furthermore, the fog was coming in so neither plane could see the other.

There were time pressures and economic pressures acting together. The Pan American pilots questioned their orders to taxi on the runway, but they continued anyway. The co-pilot of the KLM flight voiced minor objections to the pilot, suggesting that they were not yet cleared for takeoff. All in all, a tragedy occurred due to a complex mixture of social pressures and logical explaining away of discrepant observations.

The Air Florida flight from National Airport, Washington, D.C., crashed at takeoff into the 14th Street bridge over the Potomac River, killing seventy-eight people, including four who were on the bridge. The plane should not have taken off because there was ice on the wings, but it had already been delayed over an hour and a half; this and other factors "may have predisposed the crew to hurry." The accident occurred despite the first officer's (the co-pilot's) concern: "Although the first officer expressed concern that something 'was not right' to the captain four times during the takeoff, the captain took no action to reject the takeoff." Again we see social pressures coupled with time and economic forces.[15]

Designing for Error

Error is often thought of as something to be avoided or something done by unskilled or unmotivated people. But everyone makes errors. Designers make the mistake of not taking error into account. Inadvertently, they can make it easy to err and difficult or impossible to discover error or to recover from it. Consider the London stock market story that opened this chapter. The system was poorly designed. It should not be possible for one person, with one simple error, to cause such widespread damage. Here is what designers should do:

1. Understand the causes of error and design to minimize those causes.
2. Make it possible to reverse actions—to "undo" them—or make it harder to do what cannot be reversed.
3. Make it easier to discover the errors that do occur, and make them easier to correct.
4. Change the attitude toward errors. Think of an object's user as attempting to do a task, getting there by imperfect approximations. Don't think of the user as making errors; think of the actions as approximations of what is desired.

When someone makes an error, there usually is good reason for it. If it was a mistake, the information available was probably incomplete or misleading. The decision was probably sensible at the time. If it was a slip, it was probably due to poor design or distraction. Errors are usually understandable and logical, once you think through their causes. Don't punish the person for making errors. Don't take offense. But most of all, don't ignore it. Try to design the system to allow for errors. Realize that normal behavior isn't always accurate. Design so that errors are easy to discover and corrections are possible.

HOW TO DEAL WITH ERROR—AND HOW NOT TO

Consider the error of locking your keys into your car. Some cars have made this error much less likely. You simply can't lock the doors (not easily, anyway) except by using the key. So you're pretty much forced