8

The Function of the Subreality Machine

Introduction

In the last chapter we showed that our unconscious mental activity places our conscious mental activity in an Information-Limited Subreality. A “subreality machine” exists in each of our brains, creating everything that we consciously experience. This is a general description of *what* is going on. In this chapter we turn our attention to the question of *why* the brain operates in this way. Science understands the human body as a collection of individual parts, with each part carrying out a specific function for the benefit of the whole. For us to understand why the brain contains a subreality machine, we need to understand the function being performed by this mental architecture.

We will look at this issue in two different ways. In the first, we examine the basic components of the subreality machine, the information processing upon which it is based. Human color perception provides the platform for us to conduct this examination. In our second approach, we investigate the specific function carried out by the subreality machine in the human brain. How can the creation of an inner reality facilitate our finding food, attracting mates, or escaping enemies? Just what problem did evolution overcome by endowing humans with a subreality machine? And of all the different information processing architectures that could have developed in the brain, why do humans have one that generates a seemingly detailed and elaborate inner reality? As we will show, the answers to these questions come from a single starting point: *it is difficult to analyze sensory data.*

Why is the Sun Yellow?

Science has known for over 100 years that light is a wave of electric and magnetic fields. We are all familiar with waves moving on the surface of water, where the distance from one crest to the next might be as small as a few inches, or as large as hundreds of feet. This distance is called the "wavelength," and is the most important parameter associated with a wave. The wavelength of light is very short, between about 400 and 800 nanometers (billionths of a meter). To scientists, the "color" of light is exactly the same as the "wavelength."

Now we want to explore how humans perceive color. The retina in the eye contains four different types of cells that are sensitive to incoming light. One of these four, called the *rods*, is used only in night vision and cannot distinguish color. This is why the world looks black and white in dim light. The other three receptor cells are called the blue, green, and red *cones*. Each cone contains a different pigment, causing it to be sensitive to a different wavelength of light. In particular, blue cones respond best to light at a wavelength of about 450 nanometers, green cones at about 550 nanometers, and red cones at about 580 nanometers. Of course, this is very simplified explanation of a complex topic.

The important point is that light in the physical universe can have any wavelength between about 400 and 800 nanometers. However, the eye separates this continuous range into only three channels. For instance, if we shine a light at 450 nanometers into a subject's eyes, the blue receptors will be mainly activated, resulting in action potentials passing along the blue neural pathway into the brain. Likewise, light at 550 and 580 nanometers causes the same events in the green and red nerve pathways, respectively. When a mixture of wavelengths enter the eye, as is the normal case, these three channels activate in varying amounts.

In short, the only thing that the human brain knows about color is what can be contained in these three channels. If neural signals are present on the blue channel, the subject will experience the color blue. Likewise, if the green or red channel is activated, the subject will report seeing green or red, respectively. Since blue, green, and red are the only “pure” colors that the human visual system can detect, we call these the *physiological primary colors.* All other colors that humans can experience are nothing more than a mixture of these three.

A good demonstration of this is provided by color televisions and computer monitors. If you look closely at the screen with a magnifying glass, you will see that the display is composed of a large number of small dots, each being either red, green or blue. By varying the relative intensity of these three basic colors, it is possible to generate all possible colors that the human visual system can perceive. However, it cannot generate all the possible combinations of wavelengths that exist in the physical universe.

Now we come to the interesting part, what the brain does with the color information that it receives. Suppose we conduct an experiment by displaying three different colored circles on a computer monitor. To start, we will make the three circles the primary colors, one red, one green, and one blue. We then tell a test subject the name of a color, and ask him to point to it on the display. Of course, he has no trouble doing this; any person with normal vision can easily recognize red, green, and blue.

But now we change the colors being displayed so that each is a combination of two primary colors. That is, one circle is blue and green, one is blue and red, and one is red and green. This is illustrated in Fig. 8-1. We then ask our subject to point to "blue-green." After looking for a few seconds, he points to the circle where the blue and green channels are simultaneously illuminated. When told that scientists call this color *cyan*, he shrugs his shoulders and says that blue-green is more descriptive. We find a similar result when we ask him to show us “blue-red,” a color also called *magenta*. Without difficulty, he points to the correct circle.

But now we find something very strange. When we ask the subject to indicate red-green he hesitates. After a few moments of thought he tells us that there is no such thing as “red-green”; it is something that he is totally unfamiliar with. When we show him the circle with the red and green channels illuminated, he protests that the color is *yellow*, and there is not the slightest thing about it that he perceives as *red-green*. He explains that red and green remind him of apples on a tree or Christmas decorations. "That's what red and green are," he insists. "The color you are pointing to makes me think of the sun and bananas."

This phenomenon is well known in science and medicine. While there are only three *physiological* primary colors (red, green and blue), there are four *psychological* primary colors (red, green, blue, and *yellow*). In other words, our brains transform a mixture of red and green into something that is not a mixture of anything. Yellow is perceived as a pure color, not a composite. Yellow is as different from red, green and blue, as red, green, and blue are different from each other.

To appreciate just how strong this effect is, consider the colors used in traffic lights. There are three conditions that must be indicated, stop, go, and caution. The colors we choose to represent these three conditions should be as different as possible, making it easy for drivers to distinguish between them. Given this, an obvious choice might be to use the three primary colors, red, green and blue. We can also identify an infinite number of *bad* choices. For instance, using forest green, lime green, and pea green would be a disaster, since they are so similar.

But now let's look at the colors that are universally accepted for this purpose, red for stop and green for go. So far so good; these two colors are as different as possible. But the color used for caution is yellow, which is a mixture of red and green entering the eye. If we consider physiology alone, this is the absolutely worst choice that could have been made. The caution light should catch our attention; it should alert us that the situation is different than it was before. But the sequence of colors: *green to green/red to red*, would seem to do the opposite

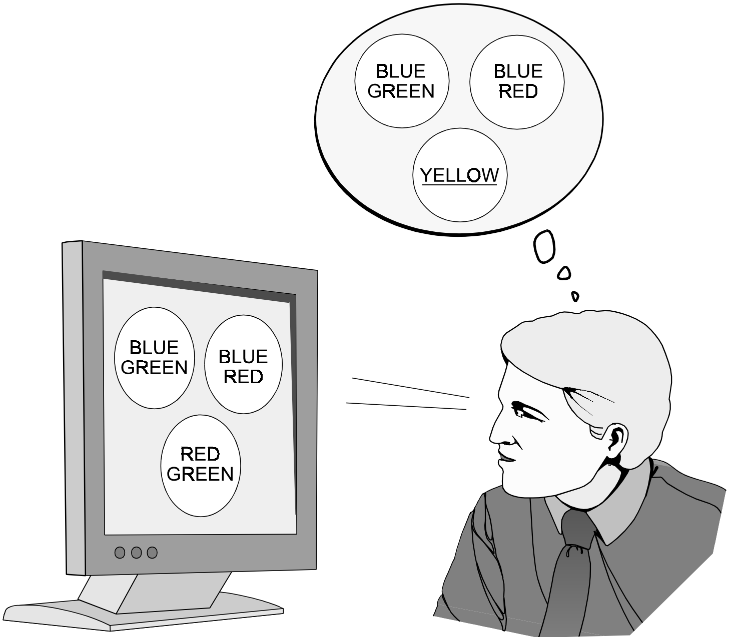

FIGURE 8-1
Color perception experiment. Humans view the combination of blue and green as a combination of blue and green. Likewise, a combination of blue and red is seen as a combination of blue and red. However, a combination of red and green is seen as *yellow*, a primary color that cannot be separated into components.

of this, minimizing the abruptness of the transitions. But, of course, it doesn't. Humans do not perceive the combination of red and green to be a combination of red and green. Rather, they perceive the combination of red and green to be yellow, a primary color in itself, something that has no relation to either red or green.

For engineers and computer scientists this is all quite uninteresting, because its explanation is so simple. As an example, suppose we asked an engineering team to create an electronic device that mimics this phenomenon. We might start with a color video camera that produces signals for red, green and blue, just as the human eye. However, the video recorder we want to use might be designed to store color from four channels, red, green, blue and yellow. The question is, how does the engineering team go about changing the data represented in three channels into data represented in four channels?

The answer is that they build a *converter*, a device that has three channels entering, and four channels exiting. The blue channel simply passes through without being altered. The other output channels (red, green, and yellow) are calculated from the other input channels (red, and green) by simple arithmetic operations, such as addition, subtraction, and comparison. Figure 8-2 shows a computer algorithm for this conversion, if you are familiar with such things. The important point is that this converter could be implemented by analog or digital electronics, computer software, a biological neural network, or any similar information processing technology. Constructing this kind of converter is extremely simple, almost trivial, to an electronic designer or computer programmer.

Now suppose we ask a scientist to examine the video recording without providing him the background on how it was made. After due inspection, the scientist proclaims that it represents a world containing *four* primary colors, red, green, blue and yellow. By this he means that each of these four colors is irreducible, and that none of these colors can be created by combining the other three. In other words, the knowledge that yellow was created from red and green is not contained within the recording. Based on the recorded video alone, yellow is as separate and distinct from red and green, as blue is from red and green.

Of course, this is exactly the situation occurring in the human visual system. Humans perceive red, green and blue as Elements-of-reality. That is, they are irreducible, they cannot be broken into more basic entities. In comparison, the colors of cyan and magenta are Information, since we perceive that they

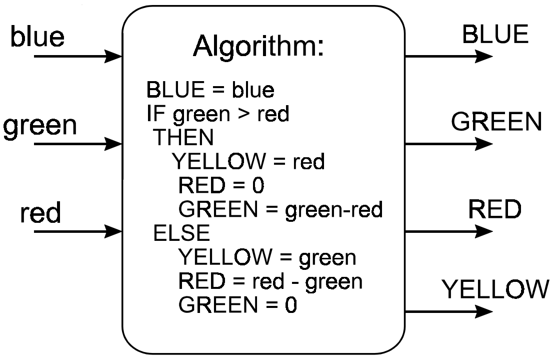

```
BLUE = blue
IF green > red
 THEN
   YELLOW = red
    RED = 0
    GREEN = green-red
 ELSE
    YELLOW = green
    RED = red - green
    GREEN = 0
```

FIGURE 8.2
Color converter. This algorithm shows how three primary colors (blue, green, and red), can be converted into four primary colors (BLUE, GREEN, RED, and YELLOW).

are composed of blue and green, and blue and red, respectively. This is just another way of saying that red, green and blue are primary colors, while cyan and magenta are not. And none of this is surprising, given that the eye inherently detects three and only three channels of color, red, green and blue.

But what about yellow? As the color signals move between the eyes and the brain, yellow is nothing more than a mixture of red and green. This means that it is *Information*, exactly the same as cyan and magenta. However, when yellow is perceived by our conscious mind, it is irreducible; it is an *Element-of-reality* of our introspective world. But as we know, nothing more than elementary operations are required to make this change, the kind of operations that are fundamental to all information processing systems. This lesson here is momentous; *the most basic operations used in information processing have the ability to change Information into Elements-of-reality.*

A critical point to understand is that changing Information into an Element-of-reality does not require that something be *added*, it requires that something be *taken away*. It is accomplished by presenting a thing, but at the same time hiding how the thing can be reduced to more fundamental components. Humans look at the color yellow and proclaim that it is irreducible, a thing in itself, an Element-of-reality. But this is a handicap, not a capability. It is a fundamental limitation on understand the thing in question. If we could look at the color yellow and perceive that it was red-green, we would be more informed, not less.

In Chapter 6 we showed that the Information-Limited Subreality has this same property, allowing the inner observer to see Elements-of-reality, while the outer observer sees only Information. We called this property the "Principle of relative reduction." This is information manipulation on a *large scale*, sufficient to manufacture an entire reality for a human or other observer. In contrast, our example of the color yellow is on a *small scale*, using the most basic information processing operations. In more poetic words, we have now examined the building and also looked at the individual bricks.

The Sensory Analysis Problem

Now we want to examine why the brain contains a subreality machine. As discussed in Chapter 3, the function of the brain is to enable *movement*, allowing the animal to locate food, escape enemies and find mates. This requires the animal to have sense organs to examine its environment, and muscles to actually move its body. The brain is the link between these two, analyzing sensory information, deciding where to move, and controlling the muscles to carry out this action. We will focus on the first of these tasks, understanding how the inner reality facilitates the analysis of sensory information. While it is possible that the inner reality is also used in determining and controlling movement, this is much more speculative and we will not pursue it here.

To start, look at the photograph in Fig. 8-3 for a few moments. When done, speak a sentence or two on what this picture is about, such as if you were briefly describing it to a friend.

FIGURE 8.3
An old photograph. This is easily recognized as a man and a woman standing in a laboratory, taken around 1900.

Your response is probably something such as: "*This is an old photograph of a middle-aged man and woman standing in a laboratory, probably taken about 1900.*" You might have even recognized it as a photograph of the great scientists Pierre and Marie Curie, famous for their work on radioactivity. You were able to extract this key information with only a few seconds of examination. It wasn't even difficult; this is a task that can be quickly carried out by any normal adult.

Now suppose that we want to build a computer to perform this same action. That is, we want to show it a picture that it has never seen before, and have it provide a short description of what the picture is about. We gather together a team of engineers and scientists that are experienced in this area, such as connecting video cameras to computers, developing software to recognize shapes in digitized images, and creating databases of stored information. We describe the goal of the project to our technical team, and ask them to give us an estimate of how long it will take, and how much it will cost. In other words, we want to get a general idea of how difficult this task really is. From a technical standpoint, is this something that is relatively easy, or is it something that is relatively hard?

After hearing our goals, most of our technical group gets up and walks out of the room, mumbling that we have wasted their time. The few that remain are kind enough to explain. One of them offers, "I rate the difficulty of new projects on a scale of 1 to 10, and this one is about 100." Another tells us, "Assuming our current rate of technological learning, this is the kind of project we might tackle 50 to 100 years from now." Still a third comments, "We have all the basic tools, but the overall complexity is just too great; it reminds me of a man holding a brick, looking up at the great pyramids."

The point is, the analysis of sensory data is *extremely* difficult, far exceeding the capabilities of present day computer technology. We perceive it as effortless only because this brain activity is blocked from our conscious examination.

The primary reason that sensory analysis is difficult rests with the data itself. The information provided by our senses is very poor quality; it is incomplete, ambiguous, contaminated with interference, and degraded in a variety of other ways. As an example, when you looked at Fig. 8-3 you probably didn't notice anything unusual. But Fig. 8-4 points out a variety of aspects of this picture that are difficult to reconcile with the physical world. For instance, some of the objects merge together without a distinct boundary between them, such as

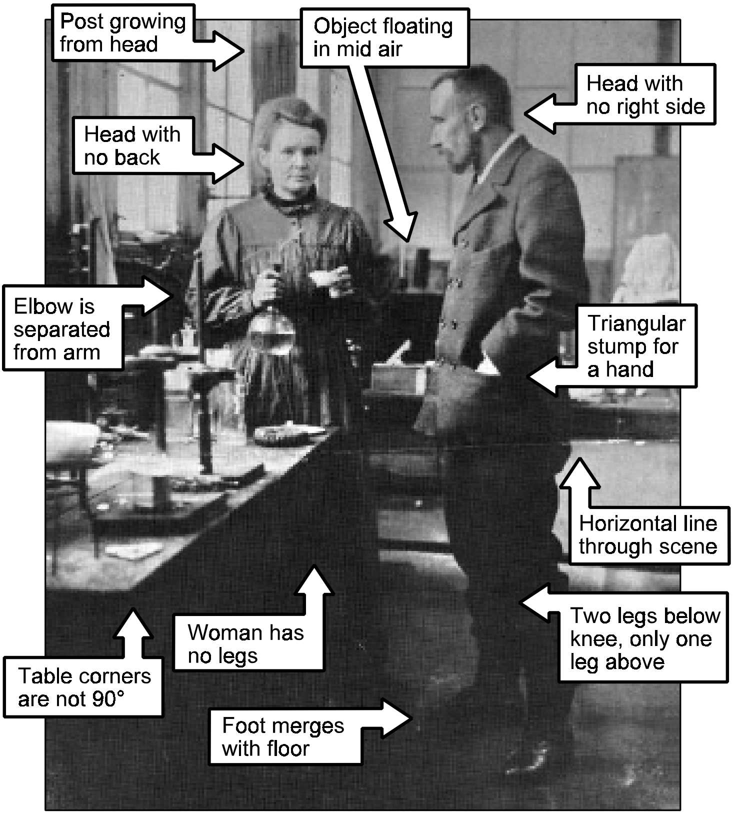

FIGURE 8-4
Image discrepancies. Vision and the other senses provide a poor representation of the physical world.

Pierre's foot and the floor. Other objects have an incomplete relationship with their surroundings, such as the dark rectangle floating in mid air. A scratch in the photograph shows up as a horizontal line, with no relation at all to the viewed scene. Severe problems are created by representing the three-dimensional setting as only a two-dimensional image. This produces missing elements, such as Marie's legs, Pierre's hand, and the back side of all the objects. It also makes whole bodies appear as discontinuous, such as the elbow being separated from the remainder of the arm. Further, the resulting geometric distortion changes the shape of objects, such as the rectangular table top appearing as a parallelogram.

Your first impression might be that the comments in Fig. 8-4 are trivial and unimportant. No so; these are problems that present day computer scientists struggle with on a day-to-day basis. But the human brain has already solved these problems; it is capable of finding the relevant data in the exceedingly poor information provided by our eyes, ears, and other sense organs. The question is, how does the brain do it so well, and what does this have to do with an inner reality?

Filtering versus Matching

To answer this question, let's look at two techniques engineers have developed to analyze poor quality data. As an example, imagine that we want to receive a radio signal from an orbiting satellite, as illustrated in Fig. 8-5. The signal being transmitted is very simple, nothing but a sine wave at a constant amplitude and frequency. This is very familiar to those who work with electronics. If you don't have such a background, just look at the pictures to get an idea of what is going on. The important point is that the signal sent by the satellite is very smooth and regular.

In an ideal situation, the signal received on the ground would be identical to the one being transmitted by the satellite. Unfortunately, this is never the case when dealing with signals that have passed through the environment. As illustrated in this figure, the received signal is very degraded; it generally resembles the transmitted signal, but it is very jagged and irregular. This is the result of many different problems. For instance, the height of the peaks may fluctuate because the satellite is in motion, or from atmospheric turbulence. In extreme cases, this can result in sections of the received signal being completely missing. Another problem is interference; for instance, our receiver might inadvertency pick up the radio

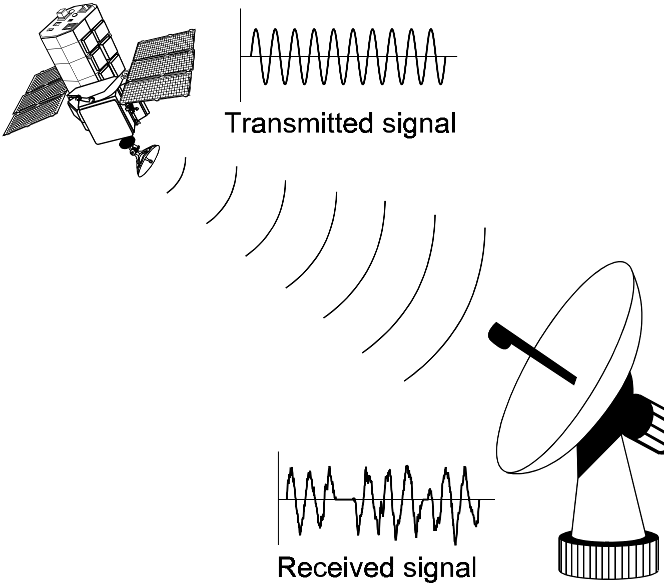

FIGURE 8-5
Passing signals through the environment. The received signal is a poor replica of the original transmitted signal, due to noise, interference, and similar problems.

transmission from an aircraft flying overhead. This becomes part of the received signal, degrading our ability to detect what is coming from the satellite. Still another problem in acquired signals is *random noise*, a term scientists and engineers use to describe a wide variety of fluctuations. This results in such things as "snow" in television pictures and static in radio broadcasts. Random noise can arise from many different sources, including the mere motion of atoms and electrons. In our example of Fig. 8-5, this type of noise shows up as a "roughness" in the received signal.

The key point is that the signal we receive on the ground is a poor quality replica of the signal transmitted by the satellite. It is distorted, missing sections, and contaminated with random noise and interference. The question is, what do we do about it? How can we change the received signal to more resemble the original?

Figure 8-6a shows our first approach to this problem, what engineers call **filtering**. There are many different ways to carry this out, and we will only give a general description leaving out the technical details. The basic idea is to pass the signal through an electronic circuit or computer routine that changes the signal's characteristics in some desirable way. For instance, if we know that the signal being transmitted from the satellite is relatively smooth, our filter might remove the roughness in the received signal, as illustrated in this figure. If you don't have a background in electronics, think of this as performing the same function as the suspension on an automobile, providing a smooth ride even over a bumpy road. Filters are very common in electronic circuits, and can be very simple to extremely complex. But even the most advanced filters have limitations on how well they can work with highly degraded data. As in this example, when interference and random noise dominate the received signal, the output of the filter still looks like interference and random noise.

Now we want to turn our attention to an alternative technique, called the **phase lock loop**. This is far less common in electronics, being used in only a few specialty applications. Just as before, we will only give a general description that leaves out the technical details. As shown in Fig. 8-6b, the phase lock loop is composed of two parts, a *comparing circuit* and a *sine wave generator*. The sine wave generator does just that; it produces a pure sine wave, without distortion, interference or noise. The function of the comparing circuit is to continually compare this created signal with the signal received from the satellite. If a difference is found between the two, the comparing circuit generates a "correction signal" that is fed into the sine wave generator. This, in turn, causes the sine wave generator to alter its output in an appropriate way to make a better match. The overall effect is that the phase lock loop generates a perfect sine wave that is the best possible match

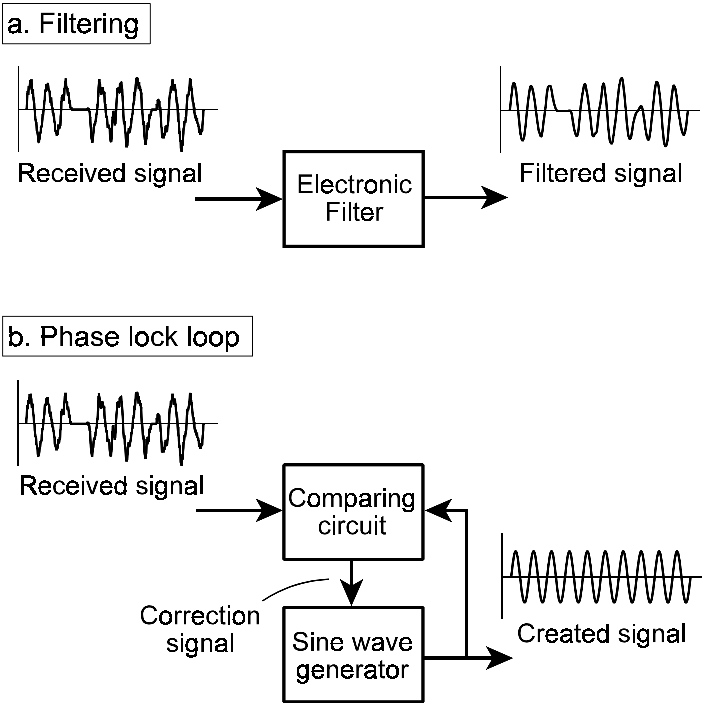

FIGURE 8-6
Filtering and PLL operation. As illustrated in (a), filtering attempts to "clean up" a contaminated signal. In comparison, (b) shows how a phase lock loop generates an entirely new signal.

to the received signal. Even if the satellite stops transmitting, the phase lock loop will still produce a pure sine wave output, its best match to the remaining random noise and interference.

The phase lock loop has one tremendous advantage and one tremendous disadvantage compared to filtering. The advantage is that it can operate with extremely high levels of interference and random noise, while still producing a near ideal output. Filtering can't come close to matching the phase lock loop in this respect. The disadvantage is that the phase lock loop only knows how to detect one very specific thing, a pure sine wave.

For instance, if the satellite started to transmit a waveform of some other shape, the phase lock loop would respond in the same old way, producing a sine wave output. In short, the phase lock loop works well with degraded data, because it is only looking for a single thing.

It is a commonplace belief that our minds directly perceive the physical universe. As an engineer would put it, the objects around us result in signals being passing into the brain, where they are somehow perceived by our conscious minds. Various filtering operations may be applied to these signals by our neural circuits, but what we end up experiencing still has a one-to-one correspondence with the external world. However, this view is simply not true. The brain does not "filter" the signals; it generates new signals that it believes are the best matches to the nearby environment. In other words, it operates like a phase lock loop, not an electronic filter.

As we move about the world in our day-to-day activities, our brains must continually keep track of what is around us. The brain is also responsible for identifying other aspects of the local environment, such as its sounds, smells, and tastes. This information about the surroundings comes to the brain through the senses, usually in a highly degraded form.

The brain's task is to extract relevant information from this jumble of interference and noise, allowing it to plan and execute movements. To do this, it takes advantage of the fact that nearly everything it encounters is *familiar*. Our daily lives are composed of objects and situations that we have experienced many times before. This means that the brain does not need to identify every possible pattern and scenario that could ever exist. On the contrary, during most of our conscious lives our brain only needs to recognize those things that it has recognized in the past. Just as the phase lock loop only looks for a single waveform, the brain only needs to look for a limited number of patterns. That is, at least most of the time.

As a demonstration of this, look at the "ambiguous" figures shown in Fig. 8-7. These are illustrations that can be interpreted

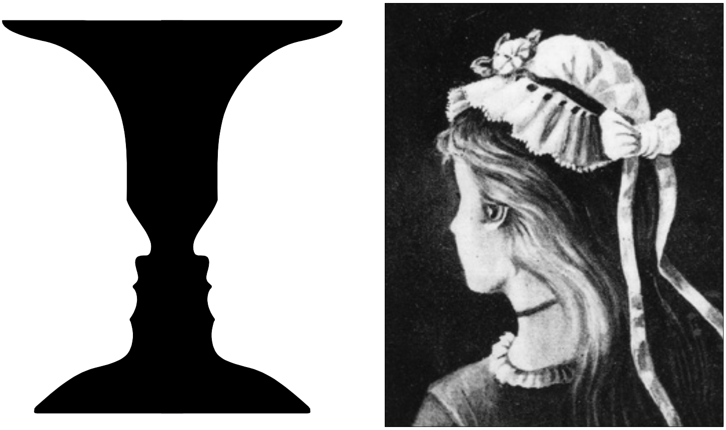

FIGURE 8-7
Ambiguous figures. On the left is “Rubin’s vase,” named after Danish psychologist Edgar Rubin who first presented it in 1921. This figure can be alternately seen as a black vase, or as two white faces in profile. The illustration on the right is often referred to as the “Boring figure,” after psychologist E.G. Boring who explored the psychology of it in the 1930s. This figure can be seen as either a young woman or an old woman. It dates to at least the 1890s, when the Anchor Buggy company used it in an advertisement with the caption: “You see my wife, but where is my mother-in-law?”

in more than one way. In (a), the image can be seen as either a black vase or two white faces. In (b), either a young woman or an old woman can be seen. However, you cannot “see” both interpretations at the same time; your mind is always locked onto one or the other. At any particular instant the figures are not ambiguous; they are a consistent representation of what you believe you are seeing. You see the vase or two faces; you see a young woman or an old woman. Even though the data entering your brain is ambiguous, your instantaneous conscious experience of the image is *not* ambiguous. Your brain has scoured the incoming data for a match. When found, you are conscious only of the consistent features of the match, not the inconsistent features of the raw data.

Let's look at an example to show just how powerful the approach of "matching" is. The images in Fig. 8-8 were created by degrading pictures of three common scenes, all of which you would immediately recognize. The resulting image quality is so poor that they hardly look like pictures at all; they seem more like random ink blots. Suppose we conduct an experiment where we show these three degraded figures to a group of 100 people and asked them to identify the pictures. How many correct responses would we expect? Of course, the answer is zero; these images are so poor that it would be impossible for anyone to do much better than guessing.

But now suppose that we redo the experiment with one significant change; we make it a multiple choice test. We start by telling our subjects that the three original images were (1) Abraham Lincoln, (2) a sunset, and (3) the Eiffel tower, in no particular order. We again ask them to identify each picture, using this additional information. After looking for a few moments, all 100 of our subjects come up with the correct answers. In other words, by narrowing the choices we have enormously improved the ability to identify patterns in ambiguous, incomplete, and noisy data. As in this example, we have changed a task that was virtually impossible, into one that can be carried out with perfect reliability.

The Subreality Machine in Operation

How does this relate to an inner reality? When we move around in the world, our brains are flooded with raw information from the senses. This data stream is so large, and such poor quality, that it would be impossible for the brain to analyze it for every possible pattern. The brain is simply not powerful enough to do this. For instance, suppose you walk into an office building for the first time. Your brain is suddenly inundated with information from your eyes and ears about the new environment. It responds by searching these data for what it expects to find, desks, chairs, people, computers, telephones, carpeting, and so on. When a match is found, the brain labels

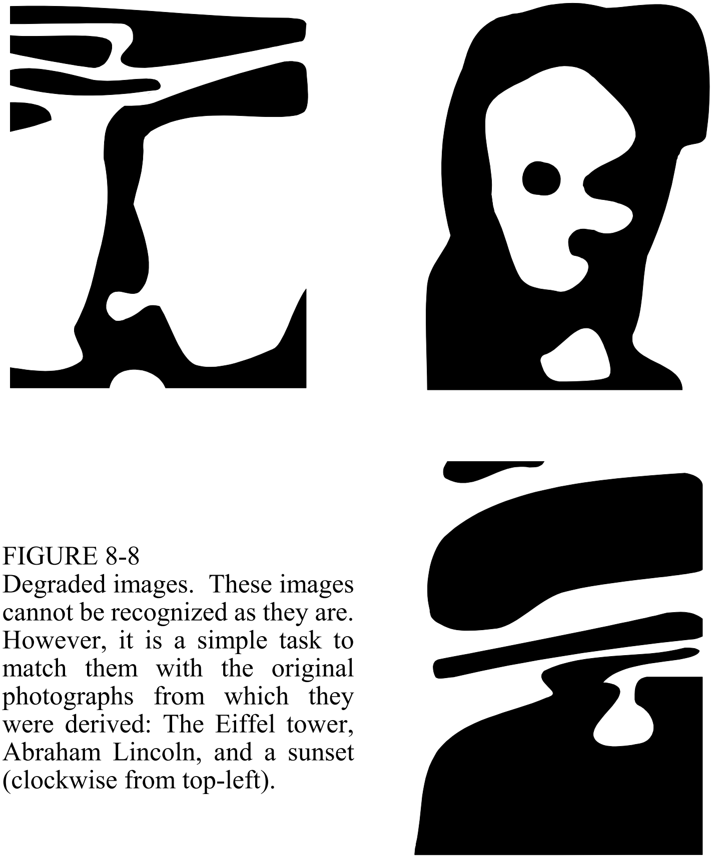

FIGURE 8-8
Degraded images. These images cannot be recognized as they are. However, it is a simple task to match them with the original photographs from which they were derived: The Eiffel tower, Abraham Lincoln, and a sunset (clockwise from top-left).

it, and then moves onto portions of the raw data that have not been recognized. This continues until the brain believes it understands the surroundings well enough to carry out its planed activities. And none of this is surprising; it is not much more than the common sense view of how our minds work.

But now let's reexamine this process using an additional assumption. We have already discussed how the analysis of sensory information is enormously difficult. Of course, this is a relative statement; it is "enormously difficult" compared to what? The assumption we will make is that sensory analysis is difficult according to two criteria, the brain's *computational power* and its *memory capabilities*.

To understand the first of these, imagine you see a chair when you walk into the new office. How long does it take you to recognize it as a chair? Of course, this happens very quickly, perhaps a tenth of a second. But how long would it take you to recognize it as one very specific chair, say, one that was part of your family's furniture when you were growing up? Since this is a more difficult task, it will take much longer, perhaps a few seconds. This is important because we live in a world where critical movements need to be made in a fraction of a second. If it took you a few seconds to identify a nearby alligator, you would be his lunch! The point is, the time it takes to complete a mental task depends on the difficulty of the task and the computational power of the brain. When we say that "sensory analysis is enormously difficult compared to the brain's computational power," we are commenting on the types of mental tasks that can be carried out within a fraction of a second. Specifically, within this key time constraint, we can sort objects into general categories, but not recognize specific entities, or search for particular characteristics.

After you enter the office and identify the chair, the next task for your brain is to take an appropriate action concerning this object. This is where the criteria concerning *memory capabilities* comes in. How do you know what this object is for, what its characteristics are, how it is used, its potential dangers, and so on? There are two obvious ways that you can obtain this information. First, your brain could search the sensory data it is receiving to answer these questions. Second, you could rely on your past experiences with this type of object. That is, you could retrieve your accumulated knowledge concerning "chairs" and assume that this particular chair has the same characteristics. Our assumption that "sensory analysis is enormously difficult compared to the brain's memory capabilities" means that the second option is faster that the first.

That is, it is faster for the brain to retrieve known information about objects in general, than it is for the brain to deduce this information each time it encounters the object.

Since the brain is a product of natural selection, it should be highly adapted to its function and environment. If sensory analysis is extremely difficult compared to the brain's computational power and its memory capabilities, this should shape the way that our mental processes are carried out. Given these assumptions, we now ask, *how would we expect the brain to operate*?

Again we will use the example of walking into a strange office. In this new situation the brain must quickly identify those things in the environment that are critical to its survival. It must do the most that it can in the first fraction of a second, the timescale that critical events happen in our world. And the best it can do is to categorize the key elements of the scene, the main features that will dictate the appropriate movements that must be made. From the sensory data, it recognizes the area as a typical office, containing a desk, chair, table, and a man. However, it determines little or nothing about the particular characteristics of these things; it only knows that they are typical members of their categories. This is all the brain can know in the first fraction of a second; its computational powers are not sufficient to extract anything else from the sensory data.

But the brain needs to have detailed information about these objects in order to move our bodies among them in a productive way. The quickest way for it to attain this information is from its own memory, what it has previously learned about objects in these particular categories. While these stored generalizations may not be accurate, they are the best that the brain can do, given the time constraint it is working under.

Keep in mind that the function of the brain can be divided into three parts, (1) analyze the sensory data to understand the environment, (2) decide where to move, and (3) coordinate the movement. Accordingly, step one must produce a "description" of the local environment that can be used by steps two and three. Given the assumptions that we have made, we would expect that this "description" would be composed of two parts, coarse information about a few key elements in the nearby environment, with the remaining details filled in from stored memories.

In short, the brain creates an inner reality that is (1) based loosely on the surroundings, (2) consistent with previous memories, and (3) free from noise, interference and ambiguity. This important concept is the sixth major teaching of the Inner Light theory:

Major Teaching #6:
The Function of the Subreality Machine

The subreality machine in the brain provides efficient sensory analysis. It achieves this by inspecting the poor quality data from the senses, and constructing an inner reality that is an estimate of the actual environment. This inner reality provides the consistent and noise-free information needed to plan and execute movements.

The Capacity of our Brains

In order for this scenario to work, the brain must have stored information about a vast number of categories of objects. This leads us to ask, is it really possible that the brain could categorize all of the familiar things that it knows? After all, we are familiar with everything from the whiskers on a cat, to the sound of a locomotive, to the taste of peanut butter. Aren't there just to many things that we are familiar with to make this possible?

To answer this question, we can make a rough estimate of just how many "things" a human knows. Of course, we can do no better than a general approximation, since we haven't defined exactly what a "thing" is. For instance, a "thing" might be the cat's whiskers, or the whole cat, or all mammals in general. Nevertheless, it is still useful to go through the calculations to get a general idea of the size of the library stored in each of our heads.

The key to making this estimate is a very simple principle: we cannot know something unless we have learned it sometime in our past. This is important, because we know very accurately how long each of us has been learning things. For instance, a typical adult has been alive for 30 years, which is the same as 10,950 days. This means they have been awake for about 175,000 hours, 10 million minutes, or 600 million seconds. The question is, on the average, how often do we learn a *new* thing? Is it every second? Every minute? Every hour?

To answer this, think about a motion picture that you saw five to ten years ago. Now suppose that you are shown a one second segment from this movie, along with a one second segment that was shot for the movie but not included in the final release. Could you reliably pick the one you had seen before? Of course not, indicating that we do not learn new things on a second-to-second time scale. But if the segments are made longer, say ten minutes, your recognition would become much more accurate. Making the segments an hour long would make your recognition nearly perfect. Using this line of reasoning, we can estimate that we learn one new "thing" about every ten minutes or so. This corresponds to about six new things per hour, 100 new things per day, 36,500 new things per year, and about *one million* new things in an entire lifetime. Keep in mind that this only pertains to long-term memory, those things that can affect our mental capabilities years after they are learned. At this instant you can probably recall hundreds of things from the last one-hour of your life. However, nearly all these will fade away, and not become a permanent part of who you are.

In short, our brains have a mental capacity of about one million "things." For comparison, this is about the same number of sentences in an encyclopedia, giving us additional reason to believe this estimate is reasonable. Of course, this number may be off by a factor of ten or more either way, especially since we have not really defined what a "thing" is. The point is, our mental world consists of a finite number of concepts that can be manipulated. Further, this finite number is not a *trillion*, or even a *billion*, but only in the neighborhood of about *one million*.

This is important because it allows us to compare our mental capacity with the physical structure of the brain. We know that the brain is composed of about 100 billion neurons, making about 100 trillion synaptic connections. In other words, the brain contains about 100,000 neurons and 100 million synaptic connections for each concept that the mind can ever process, seemingly more than sufficient to carry out the task.

Going back to our original question, is it possible that the brain has the capacity to categorize all of the things that humans know? While much of the brain's operation remains a mystery, the answer to this question seems to be a clear yes.

On a more philosophical note, this estimate of our mental capacity is a bit unsettling, especially for scientists that are accustom to dealing with very large numbers. For instance, there are about a trillion stars in our Milky Way Galaxy, and a billion trillion atoms in a single drop of water. Compared to these enormous numbers, a brain capacity of one million concepts seems quite small and almost insignificant.

Why Do We Dream?

The Inner Light Theory provides a very specific answer to the question, *What are dreams*? Each of our minds contains a subreality machine to facilitate the analysis of sensory data. Dreams result when this machinery is operated without input from the senses, resulting in an inner reality that does not correspond to the external world. Dreams are the subreality machine running amok.

This tells us what dreams are, but it does not tell us why we should have them. Why should the subreality machine activate periodically in the night without an apparent purpose? Why isn't it always shut-off during our sleep? The Inner Light Theory does not directly answer this question. However, the mental architecture described in the previous chapters does allow us to speculate on possible reasons.

To start, we will assume that nature has some reason for disconnecting the senses from the brain at night. Perhaps this is nothing more than preventing us from stumbling around in the darkness and injuring ourselves. The question then becomes, why does the subreality machine periodically activate when the sensory input is removed?

When phrased in this way, any good electrical engineer will have an immediate answer to what is going on. Manmade signal processing systems, such as those based on electronics and computers, often employ circuits to automatically adjust their sensitivity. As an example of this, consider the operation of a handheld video camera. When used to record a loud party in bright sunlight, the sound and light levels are large enough that the device can easily operate. The camera detects this and automatically reduces the sensitivity of its audio and video circuits to avoid over-driving the recording.

But now suppose that you walk into a dim room where the people are quietly talking. The camera can no longer detect the light and sound because they are below the current sensitivity level. Consequently, the recording will be nearly black and silent. However, the camera reacts to this situation by gradually increasing the sensitivity of its video and audio circuits. For instance, the camera may slowly become ten times more sensitive to light and sound over the period of a few seconds. As soon as the camera is sensitive enough to operate properly under these low-light low-sound conditions, the sensitivity stops changing and a usable recording can be made. Of course, when you walk outside the reverse process occurs; the sensitivity of the camera will gradually decrease over the first few seconds until it is appropriate for the bright and loud conditions. In short, the sensitivity of the device automatically adjusts itself to match the level of the input signals, and requires a few seconds to react to changing conditions.

This is how the automatic adjustment is suppose to work, but engineers know that many things can go wrong. For instance, during the design of the video camera an engineer had to balance the interaction of many different parameters. This includes the maximum and minimum sensitivities, how fast the camera adapts to new input levels, and the characteristics of the audio and video signals themselves.

Suppose that during the initial product design these parameters were not set properly, such as the maximum sensitivity being too high or the adaption being too quick. What would happen? When the input signals are abruptly reduced, the sensitivity of the camera will increase as expected. However, the sensitivity will overshoot and become too great, causing the recording to be a jumble of distortion and highly amplified noise. After a second or so, the camera will realize that the sensitivity is far too high, and try to correct the situation by drastically reducing it. But just as before, it overreacts, and reduces the sensitivity to a value that is far too low. This makes the recording black and completely silent. After a short time, the camera will detect this new situation and try to correct it by greatly increasing the sensitivity, starting the whole cycle over again. In the end, the recording will show brief segments of noise and distortion, separated by sections that are black and silent.

The comparison here is obvious. Dreams are an activation of the subreality machine when the input signals are taken away, with each episode occurring for about 5-10 minutes at periodic intervals of 60-90 minutes. To an electrical engineer, this sounds like oscillation of a sensitivity adjustment circuit.