



On the nature of thought processes and their relationship to the accumulation of knowledge, Part XVI—The process of making a diagnosis

Cris Anderson, M.D.¹

¹ Southern Illinois University School of Medicine, Carbondale, IL, USA

Citation: Anderson C. On the nature of thought processes and their relationship to the accumulation of knowledge, Part XVI—The process of making a diagnosis. *Dermatol Pract Conc*. 2012;2(4):12. <http://dx.doi.org/10.5826/dpc.0204a12>.

Copyright: ©2012 Anderson This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Corresponding author: Cris Anderson, M.D. Email: crisj@me.com.

Introduction

The process of making a diagnosis is a process of problem-solving during which the diagnostician notices various attributes of a patient, recognizes associations between the attributes, and applies a classification—the disease process that can then be treated, hopefully to good result for the patient. But at the heart of it, “making a diagnosis” is a very human endeavor; “to classify is human.” Our training helps us to be systematic in applying the label that we call a diagnosis to a patient, but the process of classification is something that we are born to do. Problem-solving necessarily involves a sense of uncertainty. If one knew the answer ahead of time, one would not have “a problem” to begin with.

As we all know, there is a very large amount of potential data out there and it can be difficult, especially with the time constraints we face, to figure out which attributes of the patient we should look for (what data we will need to consider). Also, although we do not like to admit it, we know our cognitive abilities are somehow limited, throwing yet another chink into the process.

Basically, this essay is an inquiry into how a finite mind can work efficiently and with purpose in what is for practical purposes a world with infinite data. We will examine various types of uncertainty and consider the implications of each type with respect to the task at hand. We will look at the various elements that make up the problem-solving process of arriving at a diagnosis and see how those elements can interact with each other, sometimes with surprising results.

Inherent in the concept of diagnosis is that illness is involved. Somehow, whatever state we are in, if we can be better, we must be ill now. So to consider the problem of making a diagnosis, we should consider the concept of health—can we identify readily “health” and “illness” and differentiate between the two?

Recognizing health and illness

All of us physicians likely think we have a pretty good idea of “health.” Webster defines “health” as “the condition of an organism or one of its parts in which it performs its vital functions normally or properly”; also “flourishing condition.” Before I go too much further, please recall the discussions in earlier essays in this series about the nature of definition, both ostensive (definition by showing) and verbal, (Interpretation) [1] and the (unavoidable) ambiguity of language (Language) [2].

Most attributes, especially attributes of disease, are not readily observed by us. Parenthetically, there is a discussion about problems of taxonomy later in this essay. Imagine you are in the waiting room of a clinic. You notice men and women, girls and boys, some sit alone, others seem to be in small groups. You notice an older gentleman sitting and talking to a middle-aged gentleman. You imagine a (healthy) son has accompanied his father to the appointment. Surely the older man is ill, since the likelihood of illness increases with age. Then a name is called. The younger man stands and follows the office staff in to see the doctor. In fact, the younger

man is ill and his father has accompanied him, providing moral support. Now imagine you are downtown shopping. Many people are on the street, going into and coming out of stores. Who has a kidney transplant? Who takes medication to treat high blood pressure? Who has high cholesterol? Most of the attributes we use to identify illness are not directly visible. Even people who have diseases of the skin generally wear clothing that prevents others from noticing a lesion.

Basically, we must decide which attribute is important to notice and how we might best notice the attribute. In medicine, the vast majority of attributes are noticed by indirect means. We order a blood test or we do a biopsy and stain the fixed tissue to highlight a potential attribute to best advantage.

In addition to the fact that attributes associated with disease may not be readily visible, members of the population, including providers of health care, do not agree on whether an attribute identified is actually within the spectrum of “health” or “disease.” Consider the following scenarios. First, when I was in medical school on a rheumatology rotation my mentor described a situation in which a patient of his, a young man, presented complaining of back pain when he lifted heavy weights as part of his exercise program. My mentor worked him up and found he had a spondylolisthesis of L4-L5. My mentor explained the situation to him, that it was a congenital condition that would require back surgery to “fix,” but that if the patient would not lift such heavy weights he would not have the back pain to begin with and would not require surgery. In fact, if the patient had not been lifting such heavy weights, he may have gotten through his life without ever knowing about his condition. The patient, however, insisted that the surgery be performed because he wanted to pursue a career in government service that required an extraordinary degree of physical conditioning. The surgery was performed and, as far as my mentor knew, the patient did well and gained employment in the job of his dreams. But was the patient ill? How many of us can perform extraordinary physical feats, say on a par with an olympic athlete? Are we unhealthy?

And what about Latisse, the glaucoma drug bimatoprost marketed as promotion for its side effect of causing longer thicker eyelashes? Are people with puny eyelashes necessarily ill?

Should every human being on the face of Planet Earth take a statin drug to lower cholesterol—this seems to be what “Big Pharma,” large pharmaceutical companies, would like us to believe.

Consider the other end of the spectrum. Ross Upshur, in “Looking for Rules in a World of Exceptions” [3], writes the following about one of his patients, “Consider the following patient. Mrs G. is 82 years old. When I assumed her care six years ago, she was given a prognosis of six months to live

from severe congestive heart failure. Mrs. G has lived beyond her original six-month prognosis. Would one consider her in good health? I don’t know. To consider her healthy is not in any way correct. To call her unhealthy is also seemingly inappropriate. I believe she is in equilibrium. [Upshur catalogs Mrs. G’s experiences with (both the diseases themselves and the reaction to her by healthcare providers) endometrial adenocarcinoma, flash pulmonary edema, aortic valve replacement, and type 2 diabetes mellitus] . . . she now requires a complex regimen of medications [including] diuretics, anti-hypertensives, cardiac medications, cholesterol-lowering medications, and diabetic medications . . . the critical determinant of why Mrs. G is still alive today . . . [has less to do with evidence-based therapy and more] to do with the fact that she has an adult son who is developmentally delayed and is absolutely dependent on her for emotional and psychological survival.”

Is a person who has never had a complaint of ill health and who feels well an hour before s/he drops dead from a heart attack healthy? The point I wish to make is that the same definition of health does not apply in every situation. A determination of health depends very much on context and perspective. The context and perspective of the patient, practitioner, and society all interweave to decide who is healthy and who is ill. Before we proceed to the process of making a diagnosis, we will start with an ideal situation to which we can compare our reality.

The ideal—by which we are able to make a correct diagnosis every time

Basically, we require three things to be able to make a diagnosis: 1) a knowledge base, 2) a set of reasoning skills, and 3) the ability to obtain specific data in the case of an individual patient. Ideally, every disease would be fully understood and specific criteria would be defined to allow diagnosis of that, and only that disease. Additionally, no diseases would share identical criteria—there would always be at least one attribute that would differentiate between similar diseases (that is to say, there would not be two or more names for the same disease). Second, the criteria used to define the disease would be unambiguous. That is to say, the same observer would make the same assessment of each attribute every time (intra-observer consistency) and different observers would all make the same assessment of the same attribute (inter-observer consistency).

Furthermore, every diagnostician would have access to the entire knowledge base and the knowledge base itself would be internally consistent (all “facts” would be satisfiable with all other “facts”—there would be no contradictions or paradoxes arising in head-to-head consideration of

different “facts”). Also, in conjunction with the two previous paragraphs, specific and parsimonious criteria would be defined for each disease and these criteria would be agreed upon by all to be sufficient to make a specific diagnosis such that disease A, and only disease A, is defined by the chosen subset of data selected from all of the data known about disease A. This item addresses the topic in information theory of “data compression” and relates to the desirability of making a diagnosis with a minimum of ancillary testing (in order to save time and money within the health care system). Additionally, the inner workings of each diagnostician would be logically consistent. Each diagnostician could assume confidently that, as each attribute of the patient is learned, all relevant data from memory would “pop” into mind and allow him/her to follow the algorithm of the process of diagnosis to *the* diagnosis, which would be correct in every case. Obviously, in the system as a whole, we are nowhere near to our ideal and, in fact, we never can be since the universe, as discussed in earlier essays in this series, is non-deterministic.

Types of uncertainty

There are multiple types of uncertainty and each type has different implications for us in our task of making a diagnosis. Two types of uncertainty are trivial. First, we might have known some fact and forgotten it. Second, perhaps we have not learned a fact yet, but some one else knows. For these two types of uncertainty, all we must do is look up the answer. A third type of uncertainty is real and unresolvable. It arises from the necessity of learning about populations by sampling. As we all know, when a new ancillary test comes out to aid in the diagnosis of a disease, it is assigned a “sensitivity” and “specificity” based on the performance of the test *in the original study population*. As we will discuss later in this essay, the attributes of the patient sitting before us do not match exactly the attributes of the study population. In fact, no two people share identical attributes (not even identical twins). To confound the issue even further, the new test is compared to a “gold standard” test—ideally a test the results of which differentiate perfectly all people who have the disease in question from all people who do not have the disease. Of course, we realize that no existing “gold standard” is ideal.

Our entire system of making diagnoses and utilizing ancillary testing is rooted in the concept that *someone knows*, somehow, who really has a disease and who does not. This is simply not true. All we have is a group of people; each person of the group has numerous attributes (some of which they share in common with other members of the group and some of which they do not share). We hope we have understood causally the disease in question well enough that a core subset of attributes is shared by group members

with the disease and that a similar, but not identical, core subset of attributes is shared by the group members without the disease. We then ask the question, “does this new test differentiate reliably between the two groups?” Can we use this new test to diagnose reliably a new patient who is not a member of the original study group? Since, by definition, we will use the new test on a patient outside the original test group, uncertainty related to the incomplete overlap of attributes necessarily introduces uncertainty into the diagnosis of every patient we see. This type of uncertainty can be lessened somewhat by ensuring that relevant attributes of the study population are well-known and that the patient on whom we are using the test actually shares the relevant attributes.

A fourth type of uncertainty is related to the philosophical concept of “vagueness.” Just where do we draw the line? In an earlier essay in this series we discussed the work of Bart Kosko, in *Fuzzy Thinking* [4], about assigning in a dichotomous manner attributes that actually occur on a continuum. In his example of apples, if we have a hundred apples and try to put them into two groups, one of red apples and one of green apples, the color of some apples will be clearly overwhelmingly red or green, but many apples will have both colors and will be more difficult to assign. Examples of this “vagueness” type of uncertainty are encountered in the practice of medicine on a daily basis. For patients with shortness of breath, one must consider whether the problem is more likely to be cardiac or pulmonary in origin. We have available in our diagnostic armamentarium B-type natriuretic peptide (BNP), which is released when the heart muscle is stretched during heart failure. Using BNP as a diagnostic aid works great, doesn’t it? If the patient is short of breath and his BNP is less than 100 units, the patient can be safely classified as a pulmonary patient. If the patient is short of breath and her BNP is more than 500 units, the patient can be safely classified as a cardiac patient. Pretty nifty! But what about the patient who is short of breath and has a BNP of 300 units (half way between 100 and 500)? In this case we simply cannot use this test to make our decision—we must search for other attributes that will help us. And similar examples abound in our daily practice. The uncertainty of “vagueness,” however, can be tamed somewhat by defining more carefully the context of the patient. All we must do is identify other attributes that alter the context of the patient and make some of the attributes we have identified more helpful—more dichotomous towards making a decision. In the new context, attribute A now argues unequivocally either for or against a diagnostic possibility. More about context later.

Another type of uncertainty is that which arises in the context of experience (or lack thereof) and is the result of making decisions based on “explicit” or “implicit” knowledge. Luchins [5] avers that “explicit” knowledge is analogous to reading written directions to perform some action,

while “implicit” knowledge is analogous to how the experience of actually performing the written instructions changes how one performs the action over time (with practice). An expert (one who has performed many times the action described in the written directions), for example, executes the written directions differently from someone who is following the instructions for the first time. The expert, via feedback gleaned while watching interim events during his/her multiple attempts performing the task, alters slightly his/her interpretation of the instructions and performs the task differently, paying particular attention to one facet or another along the way. Importantly, the results of this feedback are not usually written into a new version of the instructions. In fact, because of the ambiguity of language (as discussed in the essay in this series on Language), it is probably not even possible to write reliably implicit knowledge into written instructions. Each performer of the task learns nuances from repeated performance and over time, his/her performance improves with continued iterations (the so-called “learning curve”). This fifth type of uncertainty, then is the difference between the written instructions themselves and the unwritten “value added” to the performance resulting from experience. See below the discussion of heuristics.

Yet another source of uncertainty arises from the fact that all testing is indirect. We often know what we want to know, but we cannot look for it directly. We perform a test and from the results of that test we make inferences about what we really want to know. For example, we often want to know how well the tissues of a patient are oxygenating. Tests that are used to assess this include a hemoglobin or hematocrit to assess oxygen carrying capacity and the partial pressure of oxygen in the blood. When a patient is relatively healthy, that is to say when *most* of the patient’s physiologic systems are working well, inferences made from indirect evidence work admirably. Suppose a patient comes to the doctor’s office complaining of shortness of breath on exertion. If we find that the patient has a lower than normal hemoglobin/hematocrit, we likely assume that is the reason for the symptoms, prescribe an appropriate hematinic agent, and have the patient return in a few weeks to see if the symptoms have improved and the hemoglobin/hematocrit have returned to the normal range. But many of us have likely stood at the bedside of a gravely ill patient with advanced sepsis syndrome. The skin appears dusky or pasty or gray and the hemoglobin is a little low (not low enough to explain the clinical appearance) and the pressure of oxygen is likely adequate. We have obtained our tests to try and see how well the patient is oxygenating, but his appearance itself tells us—he is not doing well. In fact, there is *no* ancillary test that truly assesses directly how oxygen is being utilized at the cellular level.

And this problem is repeated throughout our practices, day in and day out. Radiology sees “shadows,” not tumors; levels of any analyte obtained from a blood sample test the level on a “well-mixed” sample (from a fairly large peripheral vein or artery), making it more difficult to assess a focal process. Additionally, a sample is analyzed in such a way that the analyte usually reacts in a “test system,” making the analyte more visible (perhaps with a monoclonal antibody and/or a chromogen). Even when we look histologically at structure, we see artifact induced by us—the process of biopsy wrenches the tissue from the rest of the body thereby severing its ability to receive messages that direct its function, furthermore we place it in fixative to ensure it does not deteriorate, and then we slice it thinly and stain it in a variety of ways, using chemical properties to view indirectly one facet or another. Eosinophils, for example are named after their staining properties, not after any sort of function they may have. Even our own senses “process” raw data and present it to our brains in a different format than received at the receptors of the energy we sense. Then we “recognize” the data after our brain has processed it and sent a conclusion to our consciousness somewhere. We can “work around” the uncertainty arising as a result of vagaries associated with the “indirectness” of testing by performing an additional (indirect) test that examines a different aspect of the problem, thereby obtaining “convergent” evidence. Convergent evidence comes about when results of tests looking at a problem from different perspectives all support the same hypothesis.

Another type of uncertainty is a type we can only know in hindsight—the so-called “unknown unknowns.” For this type of uncertainty we may not even know what questions to ask or how to ask a question. An example of this type of uncertainty occurred relative to subclassifying types of leukemia. Circa 1970 or so, Acute Lymphoblastic Leukemia (ALL) was known and treatable; most children fared pretty well, but a small percentage did not respond as expected to therapy. At the time, the diagnosis was made by observing cells from bone marrow or peripheral blood smeared on a glass slide stained with Wright-Giemsa stain. Shortly thereafter histochemical staining techniques were developed that could differentiate B lymphocytes from T lymphocytes. It turned out that the patients with B cell ALL responded to therapy much better than the patients with T cell ALL. Pre 1970, differences existed between B cells and T cells, but we could not tell the difference. Today we have flow cytometry and Cluster Designator markers and many subsets of lymphocytes can be detected; therapies for each subtype of leukemia or lymphoma have been developed. Yet another type of uncertainty arises as a result of the ambiguity of language. This topic is addressed more fully in earlier essays in this series on “Interpretation” and “Language.” This type of

uncertainty can be minimized by paying careful attention to the context of the situation in which words are used and by using standard definitions of the words, appropriate to the context. Context will be discussed in more detail later in this essay.

Can we draw any valid conclusions at all?

Good heavens! If everything we “know” is a derivation of something else, and an inexact datum to boot, how can we make any progress? An important point to make here is that uncertainty is not the same as randomness. While we may not be able to pinpoint exactly, and while we therefore feel uncertain, about some aspect of our work, we can be confident that the “true” answer lies somewhere between a set of limits; thus, the result is not random and completely unpredictable. The most important thing we can do is to make careful observations, check the validity of those observations with others (test inter-observer agreement), consider possible patterns and/or develop hypotheses of “causation” about the set of observations, and then test the competing hypotheses. That is to say, to evaluate observations via the process we know as science. Important factors in reasoning are context and perspective. A scientifically-minded human will construct carefully a context and consider different perspectives, reasoning through the data from each of the perspectives, trying to find a “truth.” It is also important that multiple people evaluate the same data. The importance of collective efforts is that different people will likely have different perspectives. Even if two people are considering what they think is the same perspective, the differing prior experiences of the two “thinkers” will likely lead them to a slightly different view of, and conclusion from, the data, and ensuing discussion will likely further expand the joint thinking process. If different people can confirm data or reaffirm proofs of conclusions drawn, it is more likely that the data and conclusions are “true,” that is to say conform with principles of Universal Law. It is most important that we take pains to ensure that we are describing a *consistent* system. Bronowski [6] relates that Kurt Godel and Bertrand Russell have reminded us, in a consistent system there are true things that cannot be proved (Godel’s Incompleteness Theorem), but (Russell) *in an inconsistent system one can “prove” anything!*

The root of much of the uncertainty we face arises from the situation that we are faced, as a direct result of the nature of the universe in which we find ourselves, with infinite variation around a number of common themes. Human beings are very much alike (each of us shares a large number of the attributes of “humanness”), but additionally, each of us is different in some ways from all others of the set of human beings (the entire set of attributes that describes each of us

is different from the entire set of attributes that describes each and every other member of the set “human beings”). Psoriasis, for example, has multiple presentations and features, but all features and presentations share commonalities that, when present, allow us to make the diagnosis and to have certain expectations about treatments that will be efficacious. The task we face is to recognize which subset of attributes represent the commonalities required to make a specific diagnosis among the entire set of attributes that comprise the “infinite variation.”

Human senses (sight, hearing, touch, etc) work by synthesizing many bits of data almost instantaneously; in fact, we are not even sure as individuals how we accomplish this feat and are not even aware consciously of many of the bits of energy impinging on our senses. Consider the problem of recognition of human faces by computers. Computers have not done very well, especially in the earliest attempts at programming computers to recognize faces, although progress toward recognition has been made. We humans, on the other hand, usually have little trouble recognizing faces that we have seen previously, even if the face belongs to someone we do not know well. Also, we usually have little trouble recognizing a high school classmate 20 years later at a reunion, even though many features have changed (such as sagging jowls, wrinkles, gray hair, and the like).

Parenthetically, Sacks [7] relates that a small percentage of people have prosopagnosia, the inability to recognize faces. The ability to recognize faces has been very important to humans throughout our evolution because we need to remember who has treated us fairly or unfairly so we can behave appropriately at a future encounter. Good face recognition skills serve a survival advantage.

When computers proved miserable in their first attempts, the human programmers began considering just which features humans consider important. Programs started with photographs, but it turned out that computers could not recognize a recently photographed person who was now tired, for example, because eyelids were puffier or darker than the original photo for comparison. Over time, and with adjustments to programming made possible by careful study by humans of which features are more important than others, computer face recognition has become better.

For another example that “important” features can be discerned, forensic anthropologists can draw pictures of what people may have looked like from the bone structure of a skull. Drawings of proposed current appearance of a missing adult are made from photographs of those same people as children who have been abducted. Using “important” features in the photograph of the child and enhancing those features using changes expected with growth and development and aging, a drawing is created to see if anyone has

seen recently a teenager or adult who may have been the abductee.

We seem to “know” inherently what features of human faces are important, but we may not be able to articulate what those features are. Our ability to recognize faces represents a heuristic—a “short cut” we use to make decisions.

Much study has been done about heuristics. Are they intuition that can never be defined? If someone says they have a “gut feeling” or instinct, should we trust them? I think it likely that what we call instincts or heuristics can be ultimately defined, if we deem them important enough to study and solve. An elegant example is given by Gerd Gigerenzer in *Gut Feelings* [8]. He describes the “gaze heuristic.” A friend of his played baseball and was very good at catching fly balls. The player’s coach thought he was lazy because he would sometimes just trot slowly to catch the ball, and the coach thought he should run as fast as he could to where he had calculated the ball would land. When the player did this, he missed more balls than if he used his usual technique. An assumption, prior to discovering the heuristic, was that players made complex calculations about trajectories. Gigerenzer quotes Richard Dawkins from *The Selfish Gene*, “When a man throws a ball high in the air and catches it again, he behaves as if he had solved a set of differential equations in predicting the trajectory of the ball. He may neither know nor care what a differential equation is, but this does not affect his skill with the ball. At some subconscious level, something functionally equivalent to the mathematical calculations is going on.” Gigerenzer then describes the difficulty of computing a ball’s trajectory. One must consider a parabolic curve and consider air resistance and wind and initial velocity and projection angle . . . Parenthetically, I remember when taking physics courses it took me several minutes to do my calculations, even with a slide rule. Gigerenzer describes that studies were performed to see what ball players actually do when they position themselves to catch a ball and discovered a technique called the “gaze heuristic.” The “gaze heuristic” works when the ball is high in the air. The player looks at the moving ball and decides if the angle of gaze is constant or changing. The heuristic states, “Fix your gaze on the ball, start running, and adjust your running speed so that the angle of gaze remains constant.” If players do this, they do not need to consider wind speed or spin of the ball or any other variables. Gigerenzer mentions that most ball players are unaware of how they actually catch balls, and their lack of awareness does not matter if the player is successful. But importantly, Gigerenzer avers that once the mechanism of a heuristic is known, it can be taught to people less successful and improve their performance. He describes another and similar heuristic used by airplane pilots. “If another airplane approaches and you fear collision, look for a scratch on the windshield and observe whether the plane moves rela-

tive to the scratch. If it does not, dive away immediately.” Obviously, one does not want to “catch” an airplane, as one does want to catch a fly ball.

Gigerenzer further avers that “a simple rule is less prone to estimation and calculation error and is intuitively transparent.” That is to say, it is preferable to use heuristics in many situations. However, one should recognize the heuristic as a heuristic and know the mechanism by which it works. Daniel Kahneman, in *Thinking Fast and Slow* [9], quotes Herbert Simon, a long-time researcher in the psychology of accurate intuition. Speaking of Simon’s studies of chess masters, Kahneman observes, “The psychology of accurate intuition involves no magic . . . You can feel Simon’s impatience with the mythologizing of expert intuition when he writes: ‘The situation has provided a cue; this cue has given the expert access to information stored in memory, and the information provides the answer. Intuition is nothing more and nothing less than recognition.’ . . . Valid intuitions develop when experts have learned to recognize familiar elements in a new situation and to act in a manner that is appropriate to it.” Thus, expert and accurate intuition is the ability to recognize common themes in the new situation arising in an instance of (infinite) variation. And we must study the behaviors of experts to learn the mechanisms of their heuristics so that they can be shared, thereby, in keeping with the topic of this essay, improving the overall performance of diagnosticians in the healthcare system. If we want to understand how experts think, we must understand how the human mind works.

The workings of the human mind

Basically, we can only think about what pops into our minds. Furthermore, there is a limit to how many items we can ponder simultaneously. What pops into our minds depends on association—that is to say, when we try to recall something, we try to draw an association with something else that helps us remember that item. Items that pop into our minds are likely to arrive there by “similarity matching” (“it looks like . . .”), “frequency gambling” (“I have seen that often lately”), and “recency” (“I just saw that”). That item we are considering reminds us of something else, so we consider whether the new item belongs to the same class that we are reminded of. We have seen a number of one class in particular, so if we see something that shares (reminds us of) an attribute with something we see often, surely this new item is also of the frequently noted class. Gary Marcus, in *Kluge*, [10] reminds us that we have evolved the processes of thought that we now possess, so those processes must be good enough to enable us to survive long enough to reproduce others of our kind. We can learn and memorize various facts, but we must

be able to recall items when we need to. Whether we recall what we need often depends on the context of the situation in which we are trying to think and how closely that context relates to the one in which we learned the fact we now need to recall. Kahneman describes the large amount of work he and many other students of human cognitive neuroscience have performed over the past few decades and relates the conclusions they have drawn. We have covered some of the work in earlier essays in this series. Kahneman uses the metaphor of two systems of thought, which he calls System One and System Two. Kahneman observes “System 1 operates automatically and quickly, with little or no effort and no sense of voluntary control [while] System 2 allocates attention to the effortful mental activities that demand it, including complex computations. The operations of System 2 are often associated with the subjective experience of agency, choice, and concentration.” System One is actually the prime mover and is responsible for letting thoughts arrive in our conscious awareness.

Kahneman points out that we, ourselves, identify with System 2, “the conscious reasoning self that has beliefs, makes choices, and decides what to think about and what to do. Although System 2 believes itself to be where the action is, the automatic System 1 is the hero . . . The automatic operations of System 1 generate surprisingly complex patterns of ideas, but only the slower System 2 can construct thoughts in an orderly series of steps.”

Kahneman lists, in order of complexity, some examples of activities thought to be performed by System 1:

- Detect that an object is more distant than another
- Orient to the source of a sudden sound
- Complete a common phrase, such as “bread and . . .”
- Respond to a horrible picture by making a “disgust face”
- Detect hostility in a voice
- Answer $2 + 2 = ?$
- Read words on large billboards
- Drive a car on an empty road
- Find a strong move in chess, if you are a chess master
- Understand simple sentences
- Recognize that a “meek and tidy soul with a passion for detail” [discussed earlier in Kahneman’s book as a stereotype for a librarian] resembles an occupational stereotype.

Says Kahneman “We are born to perceive the world around us, recognize objects, orient attention, avoid losses, and fear spiders. Other mental activities become fast and automatic through prolonged practice.”

Kahneman describes further System Two, “The highly diverse operations of System 2 have one feature in common: they require attention and are disrupted when attention is drawn away. Here are some examples:

- Brace for the starter gun in a race
- Focus attention on the clowns in the circus

- Focus on the voice of a particular person in a crowded and noisy room
- Look for a woman with white hair
- Search memory to identify a surprising sound
- Maintain a faster walking speed than is natural for you
- Monitor the appropriateness of your behavior in a social situation
- Count the occurrences of the letter *a* in a page of text
- Tell someone your phone number
- Park in a narrow space (for most people except the garage attendant [who can use System One for this])
- Compare two washing machines for overall value
- Fill out a tax form
- Check the validity of a complex logical argument

Kahneman observes that everyone has a degree of awareness that his/her capacity to pay attention is limited. He describes the study from the book by Chabris and Simons, *The Invisible Gorilla*, in which students, directed to count passes on a basketball court were concentrating so hard that they failed to see a gorilla walk onto the court. Kahneman notes that “intense focussing on a task can make people effectively blind, even to stimuli that normally attract attention.”

Kahneman concludes that System One generates suggestions to System Two. These suggestions are impressions, intuitions, intentions, and feelings. But once endorsed by System Two, “impressions and intuitions turn into beliefs, and impulses turn into voluntary actions . . . [most of the time, and when life runs smoothly] System 2 adopts the suggestions of System 1 with little or no modification . . . [but] when System 1 runs into difficulty [for example when asked to multiply 17×24], it calls on System 2 to support more detailed and specific processing that may solve the problem of the moment . . . System 2 is [also] activated when an event is detected that violates the model of the world that System 1 maintains.”

Kahneman continues, “The division of labor between System 1 and System 2 is highly efficient: it minimizes effort and optimizes performance. The arrangement works well most of the time because System 1 is generally very good at what it does: its models of familiar situations are often accurate, its short-term predictions are usually accurate as well, and its initial reactions to challenges are swift and generally appropriate. System 1 has biases, however, systematic errors that it is prone to make in specified circumstances . . . it sometimes answers easier questions than the one it was asked, and it has little understanding of logic and statistics . . . [also] it cannot be turned off.”

Another facet of our minds is that we see what we want or expect to see. In the essay in this series on Patterns, we discussed the work of Erich Harth [11]. He described what occurs in one’s brain as one is walking along a beach. One’s eye catches sight of a round, shiny object. One’s brain tries

to make it into a coin, but one's senses can save the day by imposing reality on the situation. If one concentrates one's visual apparatus on the object and compares it to one's expectation of the appearance of a coin, one can see if the edge of the coin is oriented at right angles to the circular surface (no), or if the surface is truly round (no), or whether there is some sort of etching—the head of a former president, perhaps—on the surface of it (no); by using one's senses to compare what one sees to one's expectation of what one thinks or hopes it might be, one concludes that what one really sees is a piece of shell.

Of course, one must engage System Two in order to make the determination. A frequent occurrence of a failure in comparing what one sees with what one expects is when one tries to proofread one's own work, especially shortly after one has written it. If one lays the work down for a day or two before proofreading, one is much more likely to detect the errors that are present.

Another thing that Kahneman points out is that humans have a strong desire for everything to fit together into a logical story. In fact, whenever we have available a few data items, we invent a story to make all the facts fit together. It is important to us that what ever occurs has a cause. When we learn a few data, even if those data prove later to be unrelated in a causal manner, our System One tries to relate them in a causal way. Kahneman gives an example “Fred's parents arrived late. The caterers were expected soon. Fred was angry. You know why Fred was angry, and it is not because the caterers were expected soon. In your network of associations, anger and lack of punctuality are linked as an effect and its possible cause, but there is no link between anger and the idea of expecting caterers. A coherent story was instantly constructed as you read; you immediately knew the cause of Fred's anger. Finding such causal connections is part of understanding a story and is an automatic operation of System 1. System 2, your conscious self, was offered the causal interpretation and accepted it.”

Kahneman also points out that System 2 can only work on one problem at a time. Recall from above that maintaining a higher than usual walking speed is a System 2 activity. Kahneman performed some studies in which, while walking at a fast pace with a test subject, the test subject would be asked to perform a complex multiplication task, such as multiply 17 by 24. Each time, the test subject would slow down to complete the multiplication task. The implication of this is that humans really cannot multitask two System 2 activities at the same time. We might be able to perform a System 1 activity concurrently with a System 2 activity, but not two activities that each require our attention (a defining aspect of a System 2 activity)

Another thing about the way humans think—we tend to think automatically somebody *must* know the answer

to any question that arises. We may admit, reluctantly, that perhaps we ourselves do not know the answer to some question, but we assume that *somebody* knows. This is natural in a way. When we are children, our parents or guardians teach us about the world around us. Whenever we do not know something, we ask them and almost always an answer is forthcoming. If they do not know, we are referred to references (dictionaries, encyclopedias and the like) and the answer is there. Even when we have difficulty finding an answer, we assume the answer must be out there somewhere. It takes a long time, but eventually, especially when we get to graduate school, we begin to learn that some questions do not have satisfactory answers. We detect inconsistencies between the “answer” to this question and the “answer” to another question. We recognize that both answers cannot be true at the same time, and we can find no ready resolution to the dilemma.

Thinking recursively

Another thing about how humans think—we think recursively. Whenever we have a thought, we tend to modify that thought by something else that pops into our minds. James Reason, in *Human Error* [12], posits that people solve problems in three basic ways: skill-based, rule-based, and knowledge-based. Skill-based is used most often and relies almost entirely on the automaticity of “System 1.” Examples include tying one's shoes, answering a telephone, brushing one's teeth, or riding a bicycle. Once we learn the activity, we hardly think about it. The action just seems to occur without much conscious thought, once we decide to initiate the action. Rule-based refers to following specific rules to an end. Algorithms are a good example of rule-based actions, and we are encouraged often to use algorithms when we practice medicine. For both skill-based and rule-based activities we are familiar with the situation and we know what to expect. The only problems with execution of these activities arise when we are distracted during skill-based actions or when we misidentify the problem and choose an inappropriate rule to execute for a rule-based problem.

We use knowledge-based techniques when we face problems that are new. We have never seen anything quite like the situation we find ourselves faced with. As a result we have to figure out what to do as we go along. At each stage, as we are trying to solve the problem, we ask ourselves, “Are we any closer to the answer?” Interestingly, we solve these problems by imagining a desired result and then trying to get to that result. We say to ourselves, “It looks a little like that other problem I had, so I'll try this maneuver that worked back then.” After that step, we reassess and decide whether we

seem to be closer to our imagined goal. If so, we continue. If not, we take a step back and try another tack.

It is the process of assessment and reassessment, using the feedback we receive from observing the status of events at each step of our progress and then deciding on the next maneuver based on the information gleaned, that is the recursive process.

Interestingly, this process that we have evolved to use is modeled in manufacturing endeavors as “Good Manufacturing Processes,” or GMP. By following GMP and checking the interim product after each step, the firm has an opportunity to make alterations and save a batch of product that, if not manufactured according to GMP, might otherwise be lost.

The point I am trying to make about thinking recursively is that even though thoughts occur to us in succession, we do not really think in a linear fashion. Algorithms, however, do tend to encourage us to think linearly. Consider driving through territory new to you in an area of town that has many signs and many potential turn offs. You have a set of directions, and you know the name of the place to which you are traveling. Assume you have a fairly good sense of direction and can tell North, South, East, and West. You know you are traveling in a northerly direction. While turning here at this street and driving through two lights, then turning at a service station, etc, you become aware that you are moving in a southeasterly direction. What do you do? You still have ten steps to follow on your set directions. Do you travel in a “linear fashion,” continuing with the set of directions until you are at the last one (analogous to linear thinking)? Or do you pull over and stop, reexamine your directions to see if you might have made an error, and possibly turn around and go back to the last site where you felt you were still traveling in a northerly direction—the sense that you were meeting your expectation (analogous to recursive thinking)?

Melanie Mitchell, in *Complexity: A Guided Tour* [13], describes her experience of writing a computer program that solves problems by making analogies. Mitchell quotes Marvin Minsky, a founder of the field of Artificial Intelligence, who said, “Easy things are hard,” referring to attempts to understand some of a human’s most basic thought processes and to replicate those processes by computer programming. Mitchell says of analogy-making, “. . . analogy-making is the ability to perceive abstract similarity between two things in the face of superficial differences. This ability pervades almost every aspect of what we call intelligence.” Mitchell quotes Henry David Thoreau, “All perception of truth is the detection of analogy.”

A basic premise underlying the project was that the strategy used in solving new problems is one of “Explore and Exploit.” One thing Mitchell realized was that all possibilities must be potentially available, but they *cannot be equally* available. For example, counterintuitive possibilities must

be potentially available, but must require a cogent reason to be considered strongly enough to warrant committing significant resources for adequate exploration of that possibility. She also realized the importance of keeping a balance between exploration and exploitation. “When promising possibilities are identified, they should be exploited at a rate and intensity related to their estimated promise, which is continually being updated. [recursive evaluation] But at all times exploration for new possibilities should continue. The problem is how to allocate limited resources—. . . be they lymphocytes, enzymes, or thoughts—to different possibilities in a dynamic way that takes new information into account as it is obtained.”

Mitchell’s goal was to write a computer program called “copycat” (because a premise of the project was that “analogy-making is a subtle form of imitation”). The goal of the program was to start with the example of two given strings of letters, similar but with an alteration, and then to give the problem of a “test” string of letters for the computer to come up with an altered string that was analogous to the example. One given alteration was “*abc* morphs to *abd*.” The test was “*mrrrjjj* morphs to ?” The goal was to use concepts possessed by the program (concepts thought to underlie human ability to form analogy) to build perceptual structures. “. . . descriptions of objects, links between objects in the same string, and correspondences between objects in different strings . . . The structures the program builds represent its understanding of the problem and allow it to formulate a solution . . . the concepts [must] be adaptable to different situations . . .”

Mitchell continues, “. . . a scheme [was proposed] for exploring uncertain environments: the ‘parallel terraced scan,’ . . . In this scheme many possibilities are explored in parallel, each being allocated resources according to feedback about its current promise, whose estimation is updated continually as new information is obtained. . . . all possibilities have the potential to be explored, but at any given time only some are actively explored, and not with equal resources. When a person (. . . or an immune system) has little information about the situation facing it, the exploration of possibilities starts out being very random, highly parallel (many possibilities being considered at once) and unfocused: there is no pressure to explore any particular possibility more strongly than any other. As more and more information is obtained, exploration gradually becomes more focused (increasing resources are concentrated on a smaller number of possibilities) and less random: possibilities that have already been identified as promising are exploited.”

Mitchell’s program has subroutines such as “Slipnet” that is a ‘network of concepts, each of which consists of a central node surrounded by potential associations and slippages’; “Workspace,” ‘in which letters composing the analogy problem reside and in which perceptual structures are built on

top of letters'; "Codelets," 'agents that continually explore possibilities for perceptual structures to build in Workspace . . . [and working in teams] . . . [using a parallel terrace scan] . . . [teams of codelets] "via competition and cooperation, gradually build up a hierarchy of structures that defines the program's 'understanding' of the situation with which it is faced'; and "Temperature," "which measures the amount of perceptual organization in the system . . . high temperature corresponds to disorganization and low temperature corresponds to a high degree of organization."

Observes Mitchell, "Via the mechanisms [of the program], Copycat avoids the Catch-22 of perception: you can't explore everything, but you don't know which possibilities are worth exploring without first exploring them. You have to be open-minded, but the territory is too vast to explore everything; you need to use probabilities in order for exploration to be fair. In Copycat's biologically inspired strategy, early on there is little information, resulting in high temperature and a high degree of randomness, with lots of parallel explorations. As more and more information is obtained and fitting concepts are found, the temperature falls, and exploration becomes more deterministic and more serial as certain concepts come to dominate. The overall result is that the system gradually changes from a mostly random, parallel, bottom-up mode of processing to a deterministic, serial, focused mode in which a coherent perception of the situation at hand is gradually discovered and gradually 'frozen in.'"

It seems to me that Mitchell's program serves as a good analogy for the process we use for making a diagnosis, or for that matter for any problem-solving activity. And because we face infinite variation around a number of common themes, we must use a "knowledge-based" approach to problem-solving more often than we would like to, even if we narrow somewhat early in the process our exploration by recognizing an attribute, or group of attributes, that seem to suggest to us a specific common theme. However, as admonished by Mitchell and her program we must still "explore" to a small degree less likely probabilities. After all, if we do not consider, however briefly, a diagnosis, we will never make that diagnosis.

Dealing with large amounts of data

How do we humans deal with large amounts of data? Is more data always better? We have already learned from Kahneman's work that our minds are lazy. We know there is no way we can learn and use efficiently vast amounts of data. We need shortcuts of some sort. A common thing that we humans seem to want is some sort of "unifying theory." We think in our heart of hearts that if we have *the* rule, or a small and easily remembered set of rules, we can figure out

anything and we will not have to memorize so much and work so hard to make progress.

Surely, it is preferable to have as much data as possible—or is it? If we consider the example of Sudoku, one strategy to solving a puzzle is to start by writing all the possible candidate numbers in the top of each empty square, then look at the puzzle, including the possibilities, and try to figure out which number goes in each square. As an avid fan of Sudoku, I actually used this technique when I first started working the puzzles. But it ended up being very confusing—there was simply too much data to consider at one time and I had more difficulty solving puzzles than I now do. I learned a few strategic "tricks" and now I only write possibilities at the top of the square when I have the square down to two or three candidates.

When we make a diagnosis, I posit that it is also possible to have too much data. It is much easier to figure out a diagnosis by performing a little strategic "legwork" first (by considering the clinical definitions of the diseases on our list of differential diagnoses for the patient) and to then order judiciously a few ancillary studies to further "flesh out" the data (determined by the clinical definitions) missing from our earlier "legwork."

Is all data "information"? It seems that data can only be considered "information" in the context of the problem *as a whole*. Data that distracts us does not help us solve the problem. That extraneous data only takes up space in our Working Memory and we waste time trying, as Kahneman warns us, to make a coherent story of all the data residing in Working Memory. As Reason would say, extraneous and distracting data serve as "Nonsigns" (as opposed to "Signs," which argue in favor of an hypothesis, or "Countersigns," which argue against an hypothesis) and, therefore, Nonsigns do not serve as "information."

The importance of context

I mentioned earlier that context is vitally important to making correct inferences about data. I recall an incident that occurred early in my pathology residency; I had been on my first surgical pathology rotation for only a couple of weeks. One case I had was a keratotic skin lesion from a middle-aged person. I was still at the stage of looking through pathology textbooks and matching pictures to make a diagnosis. Paging through Lever's *Histopathology of the Skin*, I happened upon a picture that looked very like the material on my glass slide. The caption of the picture read "Acrokeratosis verruciformis of Hopf." I carefully made a note of that and went on to the next case. When I went to sign out my cases with my attending and told him my diagnosis, he said, "Have you ever made a diagnosis of acrokeratosis verruciformis of

Hopf before?” I looked quizzically at him—surely he knew this was my second week ever of surgical pathology, and on autopsy rotations we never paid much attention (rightly or wrongly) to the deceased’s skin. I said, “No.” He responded, “Well you better make the diagnosis now because you will likely never make the diagnosis again in your entire career.” I had read that the disease was rare, but the picture did look “exactly” like what was on the glass slide representing the patient’s skin lesion. At that point in my career I had very little understanding of what sorts of information could be learned from looking at tissue through a microscope. I had no concept of how much context was necessary to make a correct diagnosis. I still burn with shame whenever I think of my early diagnostic *faux pas*. But at least I learned a lesson from that experience.

Yair Newman, in “Meaning-Making in Language and Biology” [14], points out that both language and biological systems operate in recursive-hierarchical and semantically open systems. For example, he points out that a word, by itself, can have any number of meanings. But when a word is in a sentence, the meaning of that word becomes more restricted. The meaning is further restricted when the sentence is in the additional context of a paragraph. Meaning can be further restricted by the chapter and book in which the word finds itself. In this way, a finite number of symbols (for example an alphabet) can be used to make an infinite number of messages. Additionally, while one cannot understand a word without seeing the sentence it is in, neither can one understand the sentence without understanding the words. This serves as “interaction-in-context” and exhibits hermeneutic circularity (which means that there is recursive feedback between the levels of interaction).

Newman discusses further the example of protein conformation. In the context of the cellular machinery proteins fold to a certain conformation even though some of the intermediate steps might have a higher, rather than lower, energy state (which would be unexpected to occur under the hypothesis that the protein will conform only to a lower state of energy). Of course, in the context of enzymes and catalysts, the protein is able to be temporarily in less stable intermediate conformations on the way to the final conformation.

Using the example of the immune system, Newman mentions that an agent may only serve as an antigen if the immune agents (macrophages, T cells, B cells, and cytokines) act in concert to recognize the agent as an antigen. Some people are allergic to ragweed or peanuts and other people are not, for example.

Furthermore, even second messengers such as cyclic adenosine monophosphate (cAMP), although technically freely diffusible within a cell and potentially able to interact with different subsystems (having potentially different meanings),

are generated and utilized focally, near the membrane-bound proteins with which they form a cellular subsystem of action and regulation.

In making a diagnosis, context is created by looking for multiple attributes that make up the set of a specific disease. A chief complaint can mean many conditions, but by adding attributes, gradually other conditions on the list of differential diagnoses are eliminated. A context is created such that the chief complaint comes to mean only one disease. For the example of substernal chest pain, if we add the attributes of “pressure,” radiation to the left arm, no change with breathing cycle, lasts about 20 minutes and goes away, relieved by sublingual nitroglycerin, and aggravated by exercise and exposure to cold, we can exclude the possibilities of musculoskeletal pain, pleurisy, and dissecting aortic aneurysm, but we are still left with the possibilities of cardiac angina, hypertrophic cardiomyopathy, aortic stenosis, and esophageal spasm. If we then add to the context the item of systolic ejection murmur, we narrow the problem down to aortic stenosis.

Data and information; context and perspective

I have used frequently the terms “data,” “information,” “context,” and “perspective.” How are the meanings of these terms related? A “datum” is a fact that has not yet registered in a human brain. Once a datum registers in a human brain (once a person is actually paying attention to a Datum), it becomes “information.” James Gleick, in *The Information* [15], mentions the lamentation of Heinz von Foerster during an early cybernetics conference, who complained “. . . that information theory was merely about ‘beep beeps,’ saying that only when understanding begins, in the human brain, *then* information is born—it’s not the beeps.” Thus, a datum becomes information when the human mind, using System 1, begins to associate the datum with other information/data stored in the human brain.

“Context” is the “system” of interacting data/information, being considered recursively by the thinking human. “Perspective” is a sort of “lens” through which a thinking human considers the data/information and context. Perspective can be purposely altered to a certain extent, although as mentioned in the earlier essay in this series on Interpretation some of our bedrock beliefs are so ingrained in our worldview that we are not consciously aware of how we came to believe them and we may not be able to consider some perspectives that would require considering those beliefs to be false. An example of a datum not registering in the brain of a diagnostician, and thus not becoming information, might be examining the results of a complete blood count. The diagnostician glances at the entire sheet of data, but perhaps only

registers the hemoglobin, hematocrit, total white blood cell count, and platelet count, paying no attention to the mean corpuscular volume, mean corpuscular hemoglobin, red cell distribution width, or mean platelet volume. All those data values are reported by the lab on the report, but diagnosticians may not pay attention to them in a specific patient case.

Making a diagnosis

So, if we cannot make a diagnosis under ideal conditions, how do we do it? The ground rules still apply. That is to say we still require a knowledge base, a set of reasoning skills, and the ability to acquire necessary data in the case of an individual patient.

Considering how the human mind works, we must wait until some possible disease state pops into our minds. We know from the work of Kahneman and Reason, that common diseases (frequency-matching) are likely to pop up, recently thought-about (recency-matching) diseases are likely to pop up, and diseases that we are familiar with that have similar features (similarity-matching) are likely to pop up. Since the human mind works by association, as soon as the mind of a diagnostician is stimulated by hearing a chief complaint (the usual stimulus in the healthcare setting) or by noticing something amiss (seeing what could only be a melanoma on the neck of the person standing in front of us in a check-out line at the grocery store or seeing an unresponsive person, having witnessed his recent collapse in the park) or even a combination of the two (a suspected malingerer with a complaint that seems improbable in the context of additional data), that diagnostician begins using System 1 and starts to make associations and draw data from stored memory into working memory.

In addition to the fact that we must wait for some idea to pop into our mind based on what we see is that “we notice what we notice” and nothing more. Ian Stewart, in *The Mathematics of Life* [16], discusses the work of taxonomists. States Stewart, “Taxonomists quickly learned that the most important features for classification were seldom those that immediately attracted the attention of the human observer. . . . Which characters are best suited for classifying organisms? Tigers and Zebras are both striped, but that doesn’t imply that they are closely related. In fact, tigers and zebras do not belong to the same genus, to the same family or even to the same order. Tigers are in the order *Carnivora* (carnivores), but zebras are in the order *Perissodactyla* (odd-toed hoofed animals). The two species come together only on the level of their class: both are mammals. So characters that strike the eye, like the tiger’s stripes, are often less significant than subtler ones, such as how many toes the creature possesses.”

Still explaining taxonomy, Stewart also reminds us, “One of the first steps in the development of any branch of science is to find a way to organize the wealth of observations that nature presents to us, and this is especially necessary in biology, because of the vast diversity of life.” Stewart describes the use by taxonomists of cladograms, diagrams that relate branch points and their timing and that describe shared attributes and the time during evolution that the attributes split or diverged (were no longer shared by subsequent [new] groups). Each “clade” represents an ancestral organism with all of its evolutionary descendants.

Stewart mentions that constructing a clade “involves three steps: collect data on the organisms concerned, think about suitable cladograms, and choose the best of these.” From the collected data, a set of characters are selected and the candidate organisms are assigned a value for having (1) or not having (0) the attribute. Then the data is assessed as to how many organisms have the highest percentage of attributes, how many a smaller percentage and so on. Organisms more closely related share more attributes and those sharing fewer attributes are less closely related. The data is then fed into a computer and the computer generates possible cladograms. The computer then analyzes statistically the data to see which cladogram is the best fit. Then starting with the values generated, the computer re-runs the data multiple times until there is no significant difference between the previously run cladogram and the subsequent one. The process is re-run, using different attributes. The goal says, Stewart, is to find convergent evidence, “We can be very confident if different data, analyzed by different methods, lead to similar results.”

I believe this is similar to the process of figuring out which attributes associate to define a disease. If we consider the example above of substernal chest pain, we can see that a certain number of attributes are shared by the different disease entities that make up our differential diagnosis, but as the problem is considered in the context of different data (new attributes added to the mix), some possibilities become less likely (less closely related to the definition of the disease). Also, consider over time how the understanding has evolved of various disease processes and how new attributes are added to the armamentarium in order to better classify a disease. We have iterated the process of defining a disease over decades, each new study helping to find a better definition of disease, a definition that will hopefully differentiate that disease from all others.

Of course contrary to the ideal conditions for making a diagnosis, the knowledge base of each of us is limited, the knowledge base itself has items lacking because many diseases are not fully understood, and some of our current “knowledge” will prove incorrect, perhaps because data is

missing or because we have made incorrect inferences about the data we have.

The best we can do, really, is to consider a differential diagnosis based on the presenting situation, either chief complaint or observation of some aspect of the patient. The practice of medicine is a “team sport” (not necessarily a team working at the same time and on the same patient, but collectively and over time we physicians share knowledge about a population of patients), so we had best consider a differential diagnosis that is listed by an authority for the presenting situation. An authority might be a text or consensus statement from a professional society, for example. Then, in order to figure out the most important additional data to obtain, we must understand the clinical definitions of the disease entities on our list of differential diagnoses. This is in a way analogous to using a stratagem in solving a Sudoku puzzle—instead of obtaining a large number of data, some of which might be merely distracting and not helpful, we obtain the more helpful data. The clinical definitions include the clinical features, including items of history, information obtained from physical exam maneuvers, and ancillary testing *necessary* to make a diagnosis of disease “A.”

Parenthetically a clinical definition of a disease differs from a pathophysiologic definition of the disease. The clinical definition depends on the pathophysiologic definition, but consists of items readily discernible in a clinical setting. For example, the pathophysiologic definition of myocardial infarction might be “death of myocardial fibers due to occlusion of a coronary artery by a blood clot caused by the rupture of an atherosclerotic plaque, leading to segmental loss of contractility of the heart muscle in the area supplied by the occluded artery and leading to decreased cardiac output and possibly arrhythmia, etc. Evidence of myocardial death includes features histologically of, first, contraction band necrosis and, subsequently infiltration of the necrotic area by neutrophils . . .” The clinical definition is as follows: at least two of 1) chest pain consistent with ischemic chest pain (substernal, squeezing, pressure, radiation to jaw or arm, may be accompanied by diaphoresis), 2) localized ischemic changes on electrocardiogram, consistent with blockage of a coronary artery, and 3) elevated cardiac marker (troponin and/or creatine kinase (CK)-MB.

The clinical definitions are developed over time, after many studies, including postmortem examinations on prior patients, have been performed elucidating the pathophysiologic definitions. Clinical definitions emphasize sets composed of clinical features and routine diagnostic tests, each set ideally unique to the disease in question; pathophysiologic definitions often rely on specialized testing that has not yet been approved for patient testing outside the research setting.

When considering the items on our list of differential diagnoses and when looking for data that answers the ques-

tion, “Does this patient have the features required to diagnose clinically Disease ‘A,’ Disease ‘B,’ or Disease ‘C’?” we must ask the question from the perspective of each disease on the list of differential diagnoses. We must say to ourselves, “Feature One is a sign (favors the diagnosis) for Disease ‘A,’ but a countersign (argues against the diagnosis) for Disease ‘B,’” and so forth, going through each of the data from the perspective of each disease. Then the disease that has the fewest countersigns is the most likely diagnosis. Of course, any very important countersigns may cause us to broaden our list of differential diagnoses, considering a less common disease as the cause. With the data we have, we engage in an episode of recursive thinking, going back and forth in our minds between possibilities until we finally decide on one—the diagnosis.

There are two very important points to keep in the back of our minds when we are making a diagnosis. First, disease entities do not have differential diagnoses. Disease entities share attributes and groups of attributes. Only the attributes and groups of attributes themselves have differential diagnoses.

This may seem like a nit-picking distinction, but it is crucial. Think about times you have read an article about one disease or another in the context of thinking about a particular patient. What happens when you get to the list of supposed differential diagnoses for that disease? You probably think to yourself, “I would never have considered that disease in this patient.” And then what? If you believe that diseases have differential diagnoses, just because that disease was on the list of differential diagnoses you might order another test or two to rule out the disease that you were not even considering for that patient. But if you look at it from the perspective of the previous paragraph and think about which attribute or attributes the patient has that led you to read the article to begin with, you can then decide rationally whether the disease on the list you were not considering initially should be considered (shares the attribute(s) with the disease described in the article and with your patient) or whether the disease described in the article shares *other* attributes with the disease listed as a differential diagnosis, but other attributes that your patient does not exhibit; that is to say, you should *not* consider that disease on the differential list as a possibility for your patient.

For example, if a child presents with limb pain, one possibility is sickle cell crisis. If one reads an article about sickle cell disease in all its manifestations, the list of differential diagnoses might include carotid-cavernous sinus fistula. If the patient does not have the attribute of a swollen proptotic eye, one need not consider carotid-cavernous sinus fistula as a diagnostic possibility. On the other hand, if acute leukemia is on the list of differential diagnoses, sickle cell crisis can share the attribute of limb pain with a presentation of acute

leukemia and leukemia should be considered in the differential because both conditions share the attribute exhibited by the patient—limb pain.

Disease states can be considered sets of attributes. Each disease-set has many attributes and only a subset (of the disease-set as a whole) of the attributes serve to define the disease for diagnostic purposes. For example, fever is a common symptom and is an attribute shared by many diseases, but usually we recognize that fever means “the patient is ill” and we look for other attributes that define more closely the nature of the disease. Periodicity of the patient’s fever may suggest malaria, for example, or presence of a rash accompanying the fever may suggest measles as another example; and usually we order an ancillary test, perhaps a Wright-Giemsa stained smear of peripheral blood to look for malarial parasites.

When considering the concept of disease-sets when making a diagnosis, it can be useful to consider Venn diagrams, whereby one looks for areas of overlap between sets. Looking for attributes in the area of overlap, that is to say looking for *shared* attributes, does not help us differentiate between disease-sets. We must look for attributes *outside* the area of overlap to differentiate between two (or more) diseases.

The second very important point is that there is one inviolable rule that applies to the process of making a diagnosis. That rule is Baye’s Rule. For a population in which there are members of different classifications, but with each member of the population sharing a given attribute or group of attributes, the likelihood that a specific classification should be applied to a member of that population is directly proportional to the prevalence of the classifications in the population. For example, if the population consists of 1110 items and if 1000 items are class A, 100 items are class B, and 10 items are class C, if one pulls one item “out of the hat (the hat representing the attribute or group of attributes in question)” at random, the likelihood is much higher that the item will be of class A than of any other class. Items of class A are ten times more likely than class B and 100 times more likely than class C.

As we add attributes to the set that represents the attributes of our patient, some differential diagnostic possibilities drop out of contention because those deleted diseases have attributes that exclude them from consideration or because they do not have an attribute or attributes required to make a diagnosis, leaving fewer and fewer possibilities. For a disease entity defined precisely, when the attributes of the patient match the set of attributes that make up that disease, the prevalence of the disease in a population of patients sharing all the attributes of our patient approaches 100 percent (of course there will always be some disease not yet discovered that cannot be excluded or some disease we failed to consider; thus the prevalence cannot reach 100%).

The problem, of course, in using this inviolable rule arises from trying to assign the attribute to begin with, particularly if that attribute is “a matter of degrees,” like the red and green apples or patients with shortness of breath with a measurement of B-type Natriuretic Peptide level. Attributes that lie on a continuum can only be assigned as “yes/no” in context, and by a process of recursive analysis. We must continuously “shuffle” or “juggle” the data, considering slightly new perspectives with each additional datum thrown into the mix until any change in the “apparent answer” becomes insignificant. When the degree of change has become insignificant, then we have “homed in” on a “common theme,” the insignificant change representing part of the infinite variation that is an integral part of the system in which we live and work.

Failure in making a diagnosis

Using the analogy of Explore and Exploit, we can fail at a number of places on the road to making a correct diagnosis. We can fail to recognize an attribute and not even realize a diagnosis should be made. We can “home in” on an attribute that is not important, or fail to follow up by looking for more important attributes (from the standpoint of accurate classification) and end up making a wrong diagnosis. We can look for “confirmatory” attributes only (from the perspective of only one potential diagnosis on our list of differential diagnoses), failing to rule out other items on the list of differential diagnoses. We must remember that we diagnosticians are human beings first and diagnosticians second. We think just like all other human beings. We must remember to engage System 2 to ensure we are not making a careless mistake. In the end, every human being classifies items many times a day, but as diagnosticians, our classifications affect another human life and we must take care when we perform our task that we have taken reasonable precaution to avoid the errors common to the process of classification, errors that are most often due to “signing off” too soon on the work of System 1.

Conclusion

The process of making a diagnosis is analogous to a complex adaptive system and, therefore, principles that apply to complex adaptive systems apply to the process of making a diagnosis. Complex Adaptive Systems are composed of interacting elements, each part doing a different job, but each part integral to the outcome. While we can observe and study elements, we can not observe and study interactions in the same way. We can only observe outcomes, that is to say changes to an element, that result from the interaction(s). Furthermore, no one element controls the system; however, *any one* element can affect *all other* elements. Some of the elements that

make up the system that is the process of making a diagnosis include the subsystems of the human mind (perception, interpretation, imagination, classification, attempts to construct coherence between items of information, knowledge base, ability to recall stored items into working memory, and in general all the items attributed by Kahneman to System 1 and System 2) and the relationship of each human's mind to culture and society (including our ability collectively to study disease processes and our collective understanding of the concepts of health and disease). Any one of these elements can, and does, influence the ultimate diagnosis of a given patient at a given time.

Complex systems, by definition are not completely predictable. To review from Edelman and Tononi [17] "Only something that appears to be both orderly and disorderly, regular and irregular, variant and invariant, constant and changing, stable and unstable deserves to be called complex." That interactions themselves are not observable directly explains in large part the unpredictability that occurs in complex systems. Systematic errors occur in all systems and these errors are predictable but the timing of these errors is not predictable (although systematic errors do not necessarily occur frequently—errors occur secondary to interactions that stress the system in some way). If the process of thinking is a system, there are predictable systematic errors that occur during processes of human thought, such as confirmation bias, failure to consider viable potential diagnoses (premature closure), and the like. Life itself is a complex adaptive system of which each of us is a member, and as a result we must expect rules to change. Some rules, basic rules, do not change—the mechanism of the hydrogen bond, for example. But strategic rules can change and we must be on our guard, continually looking and assessing "outcomes" to see if we expect them or whether we ourselves should change a strategic rule. It has been said we live in the Age of Information. We all think we know what that means—that there is "data, data everywhere." We think that *too much information* is a relatively new occurrence. However, Gleich describes the observations of Oxford scholar Robert Burton who wrote in the year 1621 (that is to say about 390 years ago, and before Isaac Newton, David Hume, Pierre LaPlace, Bertrand Russell and the many others who have tried to make sense of our world and to understand it), "I hear news every day, and those ordinary rumors of war, plagues . . . thefts . . . comets . . . of towns taken, cities besieged in France . . . Persia . . . daily musters and preparations . . . which these tempestuous times afford . . . so many men slain . . . stratagems and fresh alarms. A vast confusion of vows . . . edicts . . . lawsuits . . . grievances are daily brought to our ears. New books every day . . . whole catalogs of volumes of all sorts . . . controversies in philosophy, religion . . . Now come tidings of weddings . . . jubilees . . . sports . . . then again, as in a new

shifted scene, treasons . . . enormous villainies of all kinds . . . new discoveries, expeditions; now comical then tragical matters . . . today . . . officers created, to-morrow of some great men deposed, and then again of fresh honors conferred . . . one purchaseth another breaketh: he thrives, his neighbor turns bankrupt; now plenty, then again dearth and famine . . . Thus I daily hear, and such like."

Gleich continues, "Another way to speak of anxiety is in terms of the gap between information and knowledge. A barrage of data so often fails to tell us what we need to know. Knowledge, in turn, does not guarantee enlightenment or wisdom . . . It is an ancient observation, but one that seemed to bear restating when information became plentiful—particularly in a world where all bits are created equal and information is divorced from meaning."

It seems that the only way left to us to make progress and to diminish the anxiety (felt by all diagnosticians) associated with the glut of information is to leave the Age of Information behind us and to enter the Age of Context and Perspective. By entering the Age of Context and Perspective, "all bits" will no longer be equal and "information" will no longer be "divorced from meaning."

Each of the Ages of Mankind have formed the foundation of the next age. Principles learned in the Stone Age persisted in the Iron Age and on up through the ages of Agriculture and Industry. Information will not disappear if we leave the Age of Information. On the contrary, "information" can only gain meaning and actually inform us if we use that information "in context" and look at the information "from a variety of perspectives." By doing so, we can make a decision that makes the most sense in the context of what proves to be the best perspective.

As Minsky has reminded us, things that are easy for humans to do are sometimes difficult to study because they are so easy and it is not clear to us what we actually do, but by perseverance, we can make progress. A computer program (Copycat), can make an analogy, one of the most basic cognitive tasks humans do. Scientists have figured out the mechanisms of some heuristics, reliable short cuts to making correct decisions quickly. Although we may not know how or why we know something, Simon avers that a mechanism exists, if only we look for it, and the technique can be taught. We can put our minds to the problem of using context and perspective more often and more appropriately, thereby improving our efforts at diagnosis. While Uncertainty cannot be banished from our existence, we can ease the anxiety resulting from the gap between information and knowledge.

Summary

The process of making a diagnosis is a problem-solving activity carried out in the human brain by thinking recur-

sively using a strategy of “explore and exploit” under the condition of Uncertainty that arises from multiple sources. Required for success are a knowledge base, a set of reasoning skills, and the ability to obtain the appropriate data in the case of a specific patient. Context and the use judiciously of perspective are the keys to minimizing the anxiety that can overcome us in the cognitive gap between information and knowledge. Failures at diagnosis are due to systematic error and are both predictable in nature, if not in time, and avoidable potentially if one understands the process.

References

1. Anderson C. On the nature of thought processes and their relationship to the accumulation of knowledge, Part VIII: Interpretation. *Dermatopathology: Practical and Conceptual*. 2006;12(4):23.
2. Anderson C. On the nature of thought processes and their relationship to the accumulation of knowledge, Part XIV: Language—can it be used unambiguously? *Dermatopathology: Practical and Conceptual*. 2010;16(2):18.
3. Upshur R. Looking for rules in a world of exceptions, reflections on evidence-based practice. *Perspect Biol Med*. 2005;48(4):477-89.
4. Kosko B. *Fuzzy Thinking: The New Science of Fuzzy Logic*. New York: Hyperion, 1993.
5. Luchins D. Clinical expertise and the limits of explicit knowledge. *Perspect Biol Med*. 2012;55(2):283-290.
6. Bronowski J. *The Origins of Knowledge and Imagination*. New Haven: Yale University Press, 1978.
7. Sacks O. *The Mind's Eye*. New York: Vintage Books, 2010.
8. Gigerenzer G. *Gut Feelings*. New York: Penguin Books, 2007.
9. Kahneman D. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux, 2011.
10. Marcus G. *Kluge: The Haphazard Evolution of the Human Mind*. Boston: Mariner Books, 2009.
11. Harth E. *The Creative Loop: How the Brain Makes a Mind*. Reading, Massachusetts: Addison-Wesley Publishing Company, 1993.
12. Reason J. *Human Error*. New York: Cambridge University Press, 1990.
13. Mitchell M. *Complexity: A Guided Tour*. New York: Oxford University Press, 2009.
14. Newman Y. “Meaning-Making” in Language and Biology. *Perspect Biol Med*. 2005;48(3):317-27.
15. Gleich J. *The Information*. New York: Pantheon Books, 2011.
16. Stewart I. *The Mathematics of Life*. New York: Basic Books, 2011.
17. Edelman G, Tononi G. *A Universe of Consciousness: How Matter Becomes Imagination*. New York: Basic Books, 2000.