

This excerpt from

Philosophical Psychopathology.
George Graham and G. Lynn Stephens, editors.
© 1995 The MIT Press.

is provided in screen-viewable form for personal use only by members of MIT CogNet.

Unauthorized use or dissemination of this information is expressly forbidden.

If you have any questions about this material, please contact cognetadmin@cognet.mit.edu.

Chapter 1

Sensation and Perception (1981)

Fred Dretske

Information-processing models of mental activity tend to conflate perceptual and sensory phenomena on the one hand with cognitive and conceptual phenomena on the other. Perception is concerned with the pickup and delivery of information, cognition with its utilization. But these, one is told, are merely different stages in a more or less continuous information-handling process.¹ Recognition, identification, and classification (cognitive activities) occur at every phase of the perceptual process. Seeing and hearing are low-grade forms of knowing.

I think this is a confusion. It obscures the distinctive role of *sensory experience* in the entire cognitive process. In order to clarify this point, it will be necessary to examine the way information can be delivered and made available to the cognitive centers without itself qualifying for cognitive attributes—without itself having the kind of structure associated with knowledge and belief. For this purpose we must say something about the different ways information can be coded.

1 Analog and Digital Coding

It is traditional to think of the difference between an analog and a digital encoding of information as the difference between a continuous and a discrete representation of some variable property at the source. So, for example, the speedometer on an automobile constitutes an analog encoding of information about the vehicle's speed because different speeds are represented by different positions of the pointer. The position of the pointer is (more or less) continuously variable, and each of its different positions represents a different value for the quantity being represented. The light on the dashboard that registers oil pressure, on the other hand, is a digital device, since it has only two informationally relevant states (on and off). These states are discrete because there are no informationally relevant intermediate states. One could, of course, exploit the fact that lights have a variable intensity. This continuous property of the signal could be used to represent the *amount* of oil pressure: the brighter the light, the lower the

From *Knowledge and the Flow of Information* (Cambridge, MA: The MIT Press, 1981), pp. 135–153. Reprinted with the kind permission of the author and The MIT Press.

oil pressure. Used in this way the light would be functioning, in part at least, as an analog representation of the oil pressure.

The analog-digital distinction is usually used to mark a difference in the way information is carried about a variable property, magnitude, or quantity: time, speed, temperature, pressure, height, volume, weight, distance, and so on. Ordinary household thermometers are analog devices: the variable height of the mercury represents the variable temperature. The hands on a clock carry information about the time in analog form, but alarm clocks convert a preselected part of this into digital form.

I am interested, however, not in information about properties and magnitudes and the various ways this might be encoded, but in information about the instantiation of these properties and magnitudes by particular items at the source. I am interested, in other words, not in how we might encode information about temperature, but in how we might represent the *fact* that the temperature is too high, over 100°, or exactly 153°. What we want is a distinction, similar to the analog-digital distinction as it relates to the representation of properties, to mark the different way *facts* can be represented. Can we say, for example, that one structure carries the information that *s* is *F* in digital form, and another carries it in analog form?

For the purpose of marking an important difference in the way information can be encoded in a signal or structure, I propose to use the familiar terminology— analog vs. digital— in a slightly unorthodox way. The justification for extending the old terminology to cover what is basically a different distinction will appear as we proceed.

I will say that a signal (structure, event, state) carries the information that *s* is *F* in *digital* form if and only if the signal carries no additional information about *s*, no information that is not already nested in *s*'s being *F*. If the signal *does* carry additional information about *s*, information that is *not* nested in *s*'s being *F*, then I shall say that the signal carries this information in analog form. When a signal carries the information that *s* is *F* in analog form, the signal always carries more specific, more determinate, information about *s* than that it is *F*. Every signal carries information in both analog and digital form. The most specific piece of information the signal carries (about *s*) is the only piece of information it carries (about *s*) in digital form.² All other information (about *s*) is coded in analog form.

To illustrate the way this distinction applies, consider the difference between a picture and a statement. Suppose a cup has coffee in it, and we want to communicate this piece of information. If I simply *tell* you, "The cup has coffee in it," this (acoustic) signal carries the information that the cup has coffee in it in digital form. No more specific information is supplied about the cup (or the coffee) than that there is some coffee in the cup. You are not told *how much* coffee there is in the cup, how large the cup *is*, *how dark* the coffee is, what the shape and orientation of the cup are, and so on. If, on the other hand, I photograph the scene and show you the picture, the information that the cup has coffee in it is conveyed in analog

form. The picture tells you that there is some coffee in the cup by telling you, roughly, how much coffee is in the cup, the shape, size, and color of the cup, and so on.

I can say that *A* and *B* are of different size without saying how much they differ in size or which is larger, but I cannot picture *A* and *B* as being of different size without picturing one of them as larger and indicating, roughly, how much larger it is. Similarly, if a yellow ball is situated between a red and a blue ball, I can state that this is so without revealing where (on the left or on the right) the blue ball is. But if this information is to be communicated pictorially, the signal is necessarily more specific. Either the blue or the red ball must be pictured on the left. For such facts as these a picture is, of necessity, an analog representation. The corresponding statements (“*A* and *B* are of different size,” “The yellow ball is between the red and the blue balls”) are digital representations of the same facts.

As indicated, a signal carrying information in analog form will always carry some information in digital form. A sentence expressing *all* the information a signal carries will be a sentence expressing the information the signal carries in digital form (since this will be the most specific, most determinate, piece of information the signal carries). This is true of pictures as well as other analog representations. The information a picture carries in digital form can be rendered only by some enormously complex sentence, a sentence that describes every detail of the situation about which the picture carries information. To say that a picture is worth a thousand words is merely to acknowledge that, for most pictures at least, the sentence needed to express all the information contained in the picture would have to be very complex indeed. Most pictures have a wealth of detail, and a degree of specificity, that makes it all but impossible to provide even an approximate *linguistic* rendition of the information the picture carries in digital form. Typically, when we describe the information conveyed by a picture, we are describing the information the picture carries in analog form—abstracting, as it were, from its more concrete embodiment in the picture.

This is not to say that we cannot develop alternative means of encoding the information a picture carries in digital form. We could build a device (a buzzer system, say) that was activated when and only when a situation occurred at the source that was *exactly* like that depicted in the picture (the only variations permitted being those about which the picture carried no information). The buzzer, when it sounded, would then carry exactly the same information as the picture, and both structures (the one pictorial, the other not) would carry this information in digital form. Computer recognition programs that rely on whole-template matching routines approximate this type of transformation. (See Uhr 1973, chap. 2.) The incoming information is supplied in pictorial form (letters of the alphabet or geometric patterns). If there is an exact match between the input pattern and the stored template, the computer “recognizes” the pattern and labels it appropriately. The label assigned to the input pattern corresponds to our buzzer system. The output (label) carries the same information as the input pattern. The

information the picture carries in digital form has merely been physically transformed.

As everyone recognizes, however, such template-matching processes have very little to do with genuine recognition. As long as what comes out (some identificatory label) carries *all* the information contained in the input pattern, we have nothing corresponding to stimulus generalization, categorization, or classification. What we want, of course, is a computer program that will “recognize,” not just a letter A in *this* type font, in *this* orientation, and of *this* size (the only thing the stored template will *exactly* match), but the letter A in a variety of type fonts, in a variety of orientations, and a variety of different sizes. For this purpose we need something that will extract information the input pattern carries in *analog* form. We want something that will disregard irrelevant features of this particular A (irrelevant to its being an instance of the letter A) in order to respond to those particular features relevantly involved in the pattern’s being an instance of the letter A. We want, in other words, a buzzer system that is responsive to pieces of information the pictures (patterns) carry in analog form.

To understand the importance of the analog-to-digital conversion, and to appreciate its significance for the distinction between perceptual and cognitive processes, consider the following simple mechanism. A variable source is capable of assuming 100 different values. Information about this source is fed into an information-processing system. The first stage of this system contains a device that accurately registers the state of the source. The reader may think of the source as the speed of a vehicle (capable of going from 0 to 99 mph), and the first stage of our information-processing system as a speedometer capable of registering (in its mobile pointer) the vehicle’s speed. This information is then fed into a converter. The converter consists of four differently pitched tones, and a mechanism for activating these different tones. If the source is in the range 0 to 14, the lowest-pitched tone is heard. A higher-pitched tone occurs in the range 15 to 24, a still higher pitch from 25 to 49, and the highest at 50 to 99. These different ranges may be thought of as the approximate ranges in which one should be in first, second, third, and fourth gear, and the converter a device for alerting novice drivers (by the differently pitched tones) of the need to shift gears. The flow of information looks something like figure 1.1. What I have labeled the “Analog Representation” (the speedometer) carries all the information generated by the variable source. Since the source has 100 different possible states (all equally likely), the speedometer carries 6.65 bits of information about the source. It carries the information that the vehicle is going, say, 43 mph. This information is fed into a converter, and (assuming a speed of 43 mph) the third tone is activated. Since the third tone is activated when, and only when, the vehicle has a speed in the range 25 to 49, this tone carries 2 bits of information about the speed of the vehicle (a reduction of 100 equally likely possibilities to 25).

The output of this system is always less, quantitatively, than the input. Although 6.65 bits of information get in, something less than this comes out. What is gained

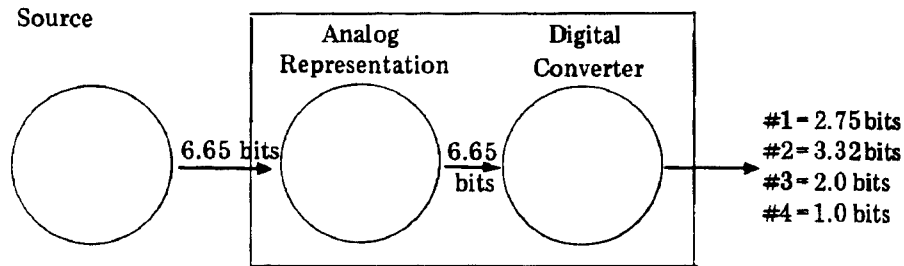


Figure 1.1

by this loss of information is a *classification* (in the various tones) of the *significant ranges* of the input variable. This is a form, albeit a fairly primitive form, of *stimulus generalization*. The output of this system ignores the difference between 43 mph and 32 mph. Both these values are treated as essentially the same. Both activate tone 3. From the point of view of the information the system is designed to communicate, the internal speedometer is an analog representation of the source because it carries more specific, more determinate information about the source than is required to control the system's output. The speedometer "says" that the vehicle is going 43 mph. Nested within this piece of information is the information that the vehicle is going *between* 25 and 50 mph. The digital converter is interested only in the latter piece of information. It "throws away" the more specific piece of information and passes along a piece of information (that the vehicle is going somewhere between 25 and 50 mph) that the speedometer carries in analog form. Of course, the speedometer carries the information that the vehicle is going 43 mph in digital form (since it carries no more specific information about the vehicle's speed), but relative to the information this system is designed to communicate (e.g., whether the speed is between 15 and 24 or between 25 and 49) the speedometer constitutes an analog representation of the state of the source. It is the information the speedometer carries in analog form that the system is *acting* on, that *drives* its motor centers (the various buzzers). The more specific pieces of information it carries are systematically ignored in order to achieve a uniform response to *relevant similarities*.

To describe a process in which a piece of information is converted from analog to digital form is to describe a process that necessarily involves the *loss* of information. Information is lost because we pass from a structure (the speedometer) of greater informational content to one of lesser information content. Digital conversion is a process in which irrelevant pieces of information are pruned away and discarded. Until information has been lost, or discarded, an information-processing system has failed to treat *different* things as essentially the *same*. It has failed to classify or categorize, failed to generalize, failed to "recognize" the input as being an instance (token) of a more general type. The simple system just described carries out this process in a completely mechanical way. Nevertheless,

although it lacks some of the essential features of a genuine perceptual-cognitive system, it illustrates the information-theoretic processes underlying all forms of stimulus generalization, classification, and recognition.

2 *Sensory vs. Cognitive Processes*

The contrast between an analog and a digital encoding of information (as just defined) is useful for distinguishing between sensory and cognitive processes. Perception is a process by means of which information is delivered within a richer matrix of information (hence in *analog* form) to the cognitive centers for their selective use. Seeing, hearing, and smelling are different ways we have of getting information about *s* to a digital-conversion unit whose function it is to extract pertinent information from the sensory representation for purposes of modifying output. It is the successful conversion of information into (appropriate³) digital form that constitutes the essence of cognitive activity. If the information that *s* is *F* is never converted from a sensory (analog) to a cognitive (digital) form, the system in question has, perhaps, seen, heard, or smelled an *s* which is *F*, but it has not *seen that* it is *F*—does not *know* that it is *F*. The traditional idea that knowledge, belief, and thought involve *concepts* while sensation (or sensory experience) does not is reflected in this coding difference. Cognitive activity is the *conceptual* mobilization of incoming information, and this conceptual treatment is fundamentally a matter of ignoring differences (as irrelevant to an underlying sameness), of going from the concrete to the abstract, of passing from the particular to the general. It is, in short, a matter of making the analog-digital transformation.

Sensation, what the ordinary man refers to as the look (sound, smell, etc.) of things, and what the psychologist refers to as the *percept* or (in some contexts) the sensory information store (SIS),⁴ is informationally profuse and specific in the way a picture is. Knowledge and belief, on the other hand, are selective and exclusive in the way a statement is. “The tapestry of awareness is rich, but the pattern recognition process, dependent on classification, is relatively impoverished in the detail with which it operates” (Pribram 1971, p. 136). Our sensory experience embodies information about a variety of details that, if carried over in toto to the cognitive centers, would require gigantically large storage and retrieval capabilities. (See Anderson and Bower 1973, p. 453.) There is more information in the sensory store than can be extracted, a limit on how much of this information can be exploited by the cognitive mechanisms.⁵

I do not mean to suggest by my comparison of sensory experience to pictures (or cognitive structures with statements) that our sensory experience is always (or *ever*) pictorial or imagistic in character—that the perception of things involves having little images (sounds, smells, tastes) in the head, or that cognitive activity is a linguistic phenomenon. It may be that the acquisition of language is essential to an organism’s having the capacity to convert sensory information into digital form (hence the capacity to have beliefs and knowledge), but this, if so, is an

empirical question, a question to which I will return in section 3 [not included herein]. For the moment I merely wish to develop the idea that the difference between our perceptual experience, the experience that constitutes our seeing and hearing things, and the knowledge (or belief) that is normally consequent upon that experience is, fundamentally, a coding difference. In this respect the relation between sensory processes and cognitive processes is like the relation between the preliminary analog representation and the subsequent digital representation described in figure 1.1. The speedometer carries the information that the vehicle is going between 25 and 50mph, and it carries this information in analog form (embedded in the more specific information that the vehicle is going 43 mph), but the particular state of the system that carries this information (the position of the pointer) is not a picture of the vehicle's speed. It does not *resemble* the state of affairs about which it carries information. And the third tone, the one that carries (in digital form) the information that the vehicle is going between 25 and 50mph, is not a *statement* or *linguistic representation* of the vehicle's speed. The conversion of information from analog to digital form *may* involve a conversion from picture to statement, but it need not. From a neurological point of view the transformation from sensory to cognitive coding takes place in the complete absence of either pictures or statements.

Unlike the simple, mechanical converter described in figure 1.1, however, living systems (most of them anyhow) are capable of modifying their digital-conversion units. As the needs, purposes, and circumstances of an organism change, it becomes necessary to alter the characteristics of the digital converter so as to exploit *more, or different*, pieces of information embedded in the sensory structures. Shifts of attention need not (although they may) involve a change in the kind of information made available in the sensory representation. There need not be any change in the way things look, sound, or smell. It may only involve a change in what pieces of information (carried in analog form) are extracted from the sensory representation.

Similarly, learning a concept is a process in which there is a more or less permanent modification of a system's ability to extract analogically coded information from the sensory store. What the simple mechanical system already described lacks is the capacity to change its response characteristics so as to exploit more, or different, pieces of information embodied in the speedometer's registration. It cannot *learn*. There is no way for it to modify the way it digitalizes information so as to respond, say, with tone 3 (or an altogether different tone) when and only when the vehicle is going between 30 and 35mph. This more specific piece of information is being picked up, processed, and fed into the converter (by the speedometer), but the system is incapable of "attending to" this fact, incapable of extracting this piece of information and "acting" on it. Contrast this with a young child, one whose receptor systems are fully matured and in normal working order, learning to recognize and identify items in her environment. Learning to recognize and identify daffodils, say, is not a process that requires the pickup of more

information from (or about) the daffodils. Given the child's keen eyesight, she may already (before learning) be receiving more information from daffodils than her more experienced, but nearsighted, teacher. Still, the teacher *knows* that the flower is a daffodil and the child does not. The child knows only that it is a flower of some sort (perhaps not even this much). What the pupil needs is not more information of the sort that could be supplied by the use of a magnifying glass. She is not *perceptually* deficient. The requisite information (requisite to identifying the flower *as* a daffodil) is getting in. What is lacking is an ability to extract this information, an ability to decode or interpret the sensory messages. What the child needs is not more information about the daffodil but a change in the way she codes the information she has been getting all along. Until this information (*viz.*, that they are daffodils) is recoded in digital form, the child *sees* daffodils but neither knows nor believes that they are daffodils.

The process of digitalization, and how it is related to learning and cognitive activity in general, will be examined at greater length in section 3. For the moment our concern is with the perceptual delivery systems—those systems whose function it is to make available, in our sensory experience, the information on which such cognitive activity depends.

It should perhaps be noted that I am greatly oversimplifying the process by means of which sensory information is extracted from the physical stimulus, integrated with collateral information, and coded in sensory form. I ignore the details of this process in order to highlight an important *difference* in the way this information is coded: a sensory (analog) form and a cognitive (digital) form. In particular, I simply ignore the fact that much of the information embodied in the sensory representation (our sensory experience) is the result of a temporal integration:

evolution has tuned the human perceptual system to register not the low-grade information in momentary retinal images but rather the high-fidelity information in *sequences of images* or in simultaneous complexes of images—the kind of information given by motion parallax and binocular parallax [Dretske's emphasis].⁶

James Gibson has argued persuasively that much of the information we manage to extract from our environment depends on a strategy of detecting higher-order invariants in a temporal series of signals—the kind of information we are able to pick up by *moving around* and registering the systematic alteration in patterns, textures, and relative positions.⁷ To understand how certain sorts of information are registered, it is important to understand the way a sensory representation may be the result of a temporal summation of signals. To think of the processing of sensory information in static terms, in terms of the kind of information embodied in the stimulus *at a particular time*, is to completely miss the extent to which our sensory representations depend on an integrative process *over time*. Even a simple

tachometer (depending, as it does, on the *frequency* of pulses) can be used to illustrate the importance of this phenomenon.

I am also ignoring the fact that our sensory representations often carry information derived from a number of different sensory channels. If we considered *only* the stimulus reaching the eyes (even if understood relative to some *temporal interval*), the conclusion would be inevitable that the stimulus is (very often at least) *equivocal*. It would be a mistake, however, to conclude from this that the sensory representation of the source is itself equivocal. For there is no reason to think that our visual experience of the source relies exclusively on the information arriving in the light reaching our *visual* receptors. Quite the contrary. Information about the gravitational orientation of objects is available in the sensory experience because the visual input is processed *jointly* with body-tilt information from proprioceptive sources. Signals specifying the position of the head in relation to gravity, the angular position and movement of the eyes in relation to the head, and the relative position and movement of all other relevant body parts play a role in determining *how* we experience what we experience. The wealth of information available in our sensory experience is to be explained, in part at least, by the fact that this experience embodies information gleaned *over time* from a *variety* of sources.

Important as it is for understanding the actual processes by means of which our sensory experience is produced, and the sorts of mechanisms responsible for the information to be found therein,⁸ the details are not directly relevant to our characterization of the result—the sensory experience itself—and the manner in which it codes information. It will be necessary, later in this chapter [not included herein], to look more closely at the machinery for delivering information in order to clarify the nature of the perceptual object and, in particular, the way the constancy mechanisms help to determine *what* we see, hear, and smell. But for present purposes these details can be set aside. Our immediate concern is with the analog character of our sensory experience.

Consider vision. You are looking at a fairly complex scene—a crowd of youngsters at play, a shelf full of books, a flag with all the stars and stripes visible. A reaction typical of such encounters, especially when they are brief, is that one has seen more than was (or perhaps *could be*) consciously noticed or attended to. There were (as it turns out) 27 children in the playground, and though you, perhaps, *saw them all*, you are unaware of how many you saw. Unless you had the time to count, you do not *believe* you saw 27 children (although you may certainly believe something less specific—e.g., that you saw *many* children or *over a dozen* children). You saw 27 children, but this information, precise numerical information, is not reflected in what you know or believe. There is no cognitive representation of this fact. To say one *saw* this many children (without realizing it) is to imply that there was *some* sensory representation of each item. The information *got in*. It was *perceptually* coded. Why else would it be true to say you saw 27 children rather than

26 or 28? Therefore, the information that is cognitively extracted from the sensory representation (the information, namely, that there are *many* children in the yard, or *over a dozen* children) is information that the sensory structure codes in *analog* form. The relationship between your *experience* of the children and your *knowledge* of the children is the same as that between the speedometer and the tone in figure 1.1.

I do not mean to be suggesting that there is a psychophysical correspondence between the information contained in the physical stimulus (or temporal sequence of stimuli) and the information contained in the sensory experience to which that stimulus gives rise. There is obviously a *loss* of information between the receptor surfaces and the internal representation. And conversely, there occurs something that is called “restoration”—an insertion into the sensory experience of representationally significant features that have no counterpart in the physical stimulus (closure of boundaries, restoration of missing sounds, etc.) (Warren 1970, p. 167). If, for example, one saw all 27 children but saw some of them only peripherally (or at dusk), it seems unlikely that information about the color of their clothes would be available in the visual experience. If such color information, contained in the stimulus (light reaching the retina), does not fall on the color-sensitive cones of the fovea, it will obviously not be available in the resulting sensory experience.⁹ But even with these peripherally seen children, information about their (rough) relative location, size, spacing, and number *will* be perceptually coded. We may suppose, along with many psychologists, that the preliminary operations associated with the preattentive processes (those which occur prior to the more elaborate perceptual processing associated with focal attention) yield only segregated figural units, units that lack the richness of information available in those portions of the visual field to which attention is given.¹⁰ Still, there is certainly more information embodied in this configuration of “figural units” than we normally extract—information about the spacing, relative size, and position of the objects represented. Typically, the sensory systems overload the information-handling capacity of our cognitive mechanisms so that not all that is given to us in perception can be digested. What is digested are bits and pieces—information the sensory structure carries in analog form.

There is a rule of seven which tells us that there is a certain limit to the rate at which subjects can process information.¹¹ When information arrives at a rate that exceeds this “capacity,” the organism fails to process it. We have already seen [“The Proper Measure of Information” of chapter 2 in Dretske 1981] that the idea of “channel capacity” has no direct application to the amount of information that can be carried by a *particular* signal. It applies only to the *average* amount of information an ensemble of signals can carry. Nevertheless, understood in the correct way, this rule seems to have some rough empirical validity. Its significance should not be misinterpreted, however. If the rule applies at all, it must be understood as applying to our capacity for *cognitively* processing information. It does not apply, and there is no evidence to suggest that it applies (quite the reverse), to our *per-*

ceptual coding of information. The rule represents some kind of limit to how much information we can extract *from* our sensory experience, not a limit to how much information can be contained *in* this experience. It assigns a limit to our capacity to convert information from analog to digital form. Recall the speedometer-buzzer system. A similar limitation applies to this system considered as a whole. Although the input contains 6.65 bits of information about the speed of the vehicle, the output contains, at most, 3.32 bits. The average output is something less than this. But this limit on the information-processing capabilities of this system is a limit that arises as a result of the analog-to-digital conversion mechanism. A full 6.65 bits of information *gets in*. There is an *internal representation* of the speed of the vehicle at all times. Nevertheless, this information is selectively utilized in order to obtain, in the output, a digital representation of certain relevant features of the input. If the rule of seven applies at all, it applies to the input-output relationship. It does not apply to that stage in the process which occurs prior to digital conversion. It does not apply to the sensory coding of information.

J. R. Pierce (1961, pp. 248–249) makes the same point in discussing the informational processing capacity of human subjects.

Now, Miller's law and the reading rate experiments have embarrassing implications. If a man gets only 27 bits of information from a picture, can we transmit by means of 27 bits of information a picture which, when flashed on a screen, will satisfactorily imitate any picture? If a man can transmit only about 40 bits of information per second as the reading rate experiments indicate, can we transmit TV or voice of satisfactory quality using only 40 bits per second? In each case I believe the answer to be no. What is wrong? What is wrong is that we have measured what gets *out* of the human being, not what goes *in*. Perhaps a human being can in some sense only notice 40 bits second worth of information, but he has a choice as to what he notices. He might, for instance, notice the girl or he might notice the dress. Perhaps he notices more, but it gets away from him before he can describe it.

Pierce is making the point that to measure the amount of information that can flow *through* a subject is to measure the limitation on the *joint* operation of the perceptual and the cognitive mechanisms (not to mention the performative mechanisms). Whatever limits are arrived at by this technique will tell us nothing about the informational limits of our sensory mechanisms. It will give us, at best, the capacity of the *weakest link* in the communication chain, and there is no reason to think that sensation constitutes the weakest link. As Pierce notes, we cannot imitate a picture with only 27 bits of information even though 27 bits of information is about the most that one can *cognitively* process. Our own perceptual experience testifies to the fact that there is more information *getting in* than we can manage to *get out*.

The same point is revealingly illustrated by a set of experiments with brief visual displays (Sperling 1960; see also Averbach and Coriell 1960 and 1961).

Subjects are exposed to an array of nine or more letters for a brief period (50 milliseconds). It is found that *after* removal of the stimulus there is a persistence of the “visual image.” Subjects report that the letters appear to be visually present and legible at the time of a tone occurring 150 milliseconds after removal of the stimulus. Neisser has dubbed this iconic memory—a temporary storage of sensory information in perceptual form (Neisser 1967, chap. 2). We need not, however, think of this as the persistence of *an image*. What persists is a structure in which incoming information *about* a pictorial array is coded in preparation for its cognitive utilization. For it turns out that although subjects can identify only three or four letters under brief exposure, *which* letters they succeed in identifying depends on the nature of a later stimulus, a stimulus that appears only 150 milliseconds after removal of the original array of letters. The later stimulus (a marker appearing in different positions) has the effect of “shifting the subject’s attention to different parts of the lingering icon.” The later stimulus changes the analog-to-digital conversion process: different pieces of information are extracted from the lingering sensory representation.

What these experiments show is that although there is a limit to the rate at which subjects can *cognitively* process information (*identify* or *recognize* letters in the stimulus array), the same limitation does not seem to apply to sensory processes by means of which this information is made available to the cognitive centers. Although the subjects could identify only three or four letters, information about *all* the letters (or at least *more* of the letters) was embodied in the persisting “icon.”¹² The sensory system has information about the character of all nine letters in the array while the subject has information about at most four. The availability of this information is demonstrated by the fact that after removal of the stimulus the subject can (depending on the nature of later stimulation) still extract information about *any* letter in the array. Hence, information about *all* the letters in the array must be available in the lingering icon. The visual system is processing and making available a quantity of information far in excess of what the subject’s cognitive mechanisms can absorb (i.e., convert to digital form). Our sensory experience is informationally rich and profuse in a way that our cognitive utilization of it is not. Relative to the information we manage to *extract* from the sensory representation (whatever beliefs may be occasioned by having this kind of sensory experience), the sensory representation itself qualifies as an *analog* representation of the source. It is this fact that makes the sensory representation more like a *picture* of, and the consequent belief a *statement about*, the source.¹³

Consider, finally, an example from developmental studies. Eleanor Gibson in reporting Klüver’s studies with monkeys describes a case in which the animals were trained to respond to the larger of two rectangles (1969, p. 284). When the rectangles were altered in size, the monkeys continued to respond to the larger of the two—whatever their absolute size happened to be. In the words of Klüver:

If a monkey reacts to stimuli which can be characterized as belonging to a large number of different dimensions, and if in doing so he reacts consistently in terms of one relation, let us say in terms of the "larger than" relation, he may be said to "abstract."

Klüver's monkeys succeeded in abstracting the larger-than relation. But how shall we describe the perceptual situation *before* they learned to abstract this relation. Did the rectangles *look* different to the monkeys? If not, how could they ever learn to distinguish between them? What possible reinforcement schedule could get them to react differently to perceptually indistinguishable elements? It seems most natural to say in a situation of this sort (and the situation is typical of learning situations in general) that prior to learning, prior to successful abstraction of the appropriate relation, the monkey's perceptual experience contained the information that it only later succeeded in extracting. It is possible, I suppose, that the rectangles only *began* to look different to the monkeys after repeated exposures, that the reinforcement schedule actually brought about a perceptual (as well as a cognitive) change.¹⁴ This would then be a remarkable case of perceptual learning (change in the *percept* or sensory representation as a result of training) (Epstein 1967). Perceptual learning may certainly take place, especially with the very young and the newly sighted, and in mature subjects with ambiguous figures¹⁵ but there is no reason to suppose that it is occurring in *every* learning situation with mature subjects. What is taking place here is very much like what takes place with the young child learning to recognize daffodils. The flowers do not look any different; the subject merely learns how to organize (recode) the information already available in its sensory experience.

The situation becomes even clearer if we present the monkeys with three rectangles and try to get them to abstract the "intermediate-size" relation. This more difficult problem proves capable of solution by chimpanzees, but the monkeys find it extremely difficult.¹⁶ Let us suppose that they are incapable of this more sophisticated type of learning. What shall we say about the perceptual situation with respect to the monkeys? Since they have already abstracted the larger-than relation, it may be assumed that they are receiving, and perceptually coding, the information that rectangle *A* is larger than *B*, and that *B* is larger than *C*. In ordinary terms this is a way of saying that the intermediate rectangle (*B*) *looks* smaller than the larger (*A*) and larger than the smaller (*C*). But information about which rectangle is intermediate, though obviously embedded (in analog form) in the perceptual experience itself, is not (and apparently cannot be) cognitively extracted by the animal. To say that the monkey cannot abstract the intermediate-size relation, therefore, is *not* to say anything about the way it perceptually codes information about figures. Rather, it is to say something about its cognitive limitations. The information is available in analog form in the experience the animal is having of the three rectangles, but the animal is unable to generate an appropriate on-off response, the kind of response characteristic of recognition or identification, to this

piece of information. It does not *know* (think, believe, judge) that *B* is of intermediate size, even though this information is available in its sensory representation of *A*, *B*, and *C*.¹⁷

Although our speedometer-tone system cannot learn, its limitations can be usefully compared with those of the monkey. This simple mechanical system can receive, process, and generate an internal (analog) representation of the fact that the vehicle is going between 30 and 35 mph. The speedometer's registration of (say) 32 mph is an analog encoding of this information. As originally conceived, however, the system as a whole cannot be made to "respond" to this piece of information. We get the same tone whether the vehicle is going between 30 and 35 mph, slower (down to 25 mph), or faster (up to 49 mph). The problem lies in the system's built-in limitation for converting information from analog to digital form. It can "recognize" a speed as between 25 and 50 mph because this fact, the fact that the speed is within this interval, is information the system is designed to convert into digital form (a distinctive tone).¹⁸ But the system is unable to "recognize" finer details, unable to make more subtle discriminations. It has no *concept* of something's being between 30 and 35 mph, no *beliefs* with this content, no internal structure with this kind of *meaning*.

To summarize, then, our perceptual experience (what we ordinarily refer to as the look, sound, and feel of things) is being identified with an information-carrying structure—a structure in which information about a source is coded in analog form and made available to something like a digital converter (more of this in section 3) for cognitive utilization. This sensory structure or representation is said to be an analog encoding of incoming information because it is always information *embedded in* this sensory structure (embedded within a richer matrix of information) that is subjected to the digitalizing processes characteristic of the cognitive mechanisms. Until information has been *extracted from* this sensory structure (digitalization), nothing corresponding to recognition, classification, identification, or judgment has occurred—nothing, that is, of any *conceptual* or *cognitive* significance.

If perception is understood as a creature's *experience* of his surroundings, then, perception itself is cognitively neutral.¹⁹ Nevertheless, although one can see (hear, etc.) an *s* which is *F* (sensorily encode information about *s* and, in particular, the information that *s* is *F*) without believing or knowing that it is *F* (without even having the concepts requisite to such beliefs), perception itself depends on there *being* a cognitive mechanism able to utilize the information contained in the sensory representation. In this sense, a system that cannot know cannot see; but if the system is capable of knowing, if it has the requisite cognitive mechanisms, then it can see without knowing.²⁰ A sensory structure that carries the information that *s* is *F* is not to be confused with a belief about *s*, a belief to the effect that *s* is *F*, but to qualify as a *sensory* representation of *s* (an experience of *s*), this structure must have a certain function within the larger information-processing enterprise. It must make this information available to a suitable converter for possible cognitive utilization. . . .

Notes

1. The following is typical: "Sensation, perception, memory and thought must be considered on a continuum of cognitive activity. They are mutually interdependent and cannot be separated except by arbitrary rules and momentary expediency." R. N. Haber, "Introduction" in Haber 1969.
2. The parenthetical "about *s*" is necessary at this point since, as we shall see in chapter 7 of Dretske 1981 [not included herein], information *about s* that is coded in digital form may nonetheless be nested in information about some other item.
3. It is not *merely* the conversion of information from analog to digital form that qualifies a system as a perceptual-cognitive system. The speedometer-buzzer system described above neither *sees* nor *knows* that the vehicle is going between 25 and 49 mph when the third tone is activated. To qualify as a genuine perceptual system, it is necessary that there *be* a digital-conversion unit in which the information can be given a cognitive embodiment, but the cognitive embodiment of information is not *simply* a matter of digitalization. What additional conditions must be satisfied to qualify a structure as a *cognitive* structure (besides digitalization) will be discussed in section 3 [not included herein, see Dretske 1981, chap. 6].
4. It has also been called the *Precategorical Acoustic Store* by R. G. Crowder and J. Morton (1969). Roberta Katzky (1975) notes that the term precategorical is important "because it implies that information held in the registers is not held there as recognized, categorized items, but in raw, sensory form . . . That the sensory registers are precategorical deserves emphasis here, because a central problem in research relating to the registers is the separation of true effects of sensory storage from possible effects of recognized information" (pp. 39–40).
5. In commenting on the SIS (sensory information storage), Lindsay and Norman (1972, p. 329) note that this "discrepancy between the amount of information held in the sensory system and the amount that can be used by later stages of analysis is very important. It implies some sort of limit on the capacity of later stages, a limit that is not shared by the sensory stages themselves."
6. Bower 1972, p. 357. Ulric Neisser also notes that the progressive deletion of microtexture at an edge yields a compelling perception of one surface going behind another and that this kind of information comes into existence only when something moves (it does not exist in the frozen array) (1977, p. 22).

In a summary of kinetic-size constancy Gunnar Johansson concludes that even under extremely impoverished stimulus conditions the sensory system is capable of extracting sufficient information (for the constancy effect) from *changing* patterns (1977, p. 382).

7. See Gibson 1966 and 1950. There may be some question of whether Gibson's notion of information is the same as that with which we are operating in this work [Dretske 1981]. In a conference on philosophy and psychology (Cornell University, April 2–4, 1976), Ulric Neisser claimed that Gibson's concept of information could be identified with Shannon's. David Hamlyn denied this, and if I understood him correctly, so did Gibson. Yet, the following passage is revealing:

Let us begin by noting that *information about* something means only *specificity* to something. Hence, when we say that information is conveyed by light, or by sound, odor, or mechanical energy, we do not mean that the source is literally conveyed as a copy or replica. The sound of a bell is not the bell and the odor of cheese is not the cheese. Similarly the perceptive projection of the faces of an object (by the reverberating flux of reflected light in a medium) is not the object itself. Nevertheless, in all these cases a property of the stimulus is univocally related to a property of the object by virtue of physical laws. This is what I mean by the conveying of environmental information. (Gibson 1966, p. 187)

This, it seems to me, fully justifies Neisser's judgment. It is, moreover, in reasonably close agreement with the concept of information developed in chapter 3 of the present work [Dretske 1981]. See Neisser 1977 and Hamlyn 1977.

8. The underlying sensory mechanisms may even involve what some investigators (following Helmholtz) are pleased to describe as *computational* or *inferential* processes. Although I see nothing

wrong with using this terminology to describe sensory processes, I think it a mistake to be (mis)led by it into assigning *cognitive* structure to such processes. We may describe sensory phenomena in informational terms, in terms that involve (to this extent at least) a structure's having a *propositional content*, but a structure's having a propositional content we associate with knowledge, belief, and judgment. I return to this point in chapter 7 [Dretske 1981].

9. Which is not to say that peripherally seen things will *look* colorless. This may be viewed as a case of perceptual restoration. The point is, however, that this restoration does not carry *information* about the color of the objects. Similarly, there is a spot on the retina (the blind spot) where the optic nerve leaves the eye which is incapable of picking up information from the stimulus. Nevertheless, if a homogeneous field (e.g., a sheet of white paper) is fixated (with one eye), we do not see a black spot. One should not suppose, however, that this sensory "interpolation" carries information about the stimulus. For, obviously, if there happened to be a black spot at this point in the field, then (under rigorously constricted viewing conditions) we would not see it. This information would be lost.
10. See, for example, Neisser 1967, pp. 94–104. Also see Hebb 1974 (pp. 140–41): "The primitive unity of a figure is defined here as referring to that unity and segregation from the background which seems to be a direct product of the pattern of sensory excitation and the inherited characteristics of the nervous system on which it acts. The unity and distinctiveness of such figures from their background, then, are independent of experience, or 'primitive.'"
11. Miller 1956. The number seven is an index to our capacity for making accurate absolute judgments of unidimensional stimuli. Our common ability to accurately identify any one of several hundred faces, any one of several thousand words, etc., should not be taken as an exception to this "rule." For faces, words, and objects are *multidimensional* stimuli.
12. "It appears as if all of the information in the retinal projection is available in this iconic storage, since the perceiver can extract whichever part is asked for" (Haber and Hershenson 1973, p. 169).
13. Irvin Rock interprets these experiments as suggesting that "in some sense of the term perception, all items in the array are perceived. Some sensory representation of each item endures for a fraction of a second. Perception during that brief period is based on the persistence in the visual system of the neural discharging triggered by the retinal image of the letters even after the letter display is turned off. Unless the items are further processed, however, these sensory representations will quickly fade away" (1975, p. 359). For the sense of the term "perception" in which all items are perceived, see "The Objects of Perception" (Dretske 1981, chapter 6, section 3).
14. But how then explain the different responses? "If experience is to have an effect, there nevertheless must first be a perception of the pattern that is itself *not* a function of experience, and through that perception the relevant memory traces can be activated on the basis of similarity" (Rock 1975, p. 361).
15. See, for example, Steinfeld 1967, pp. 505–522. Also Irvin Rock, "But there is a genuine perceptual change when in viewing potentially familiar figures one goes from an initial 'nonsense' organization to a subsequent 'meaningful' organization. The figure looks different when it is recognized" (Rock 1975, p. 348).
16. E. Gibson 1969, p. 292.
17. In his excellent introductory text, Irvin Rock (1975) is careful throughout to distinguish perceptual and cognitive issues. As a case in point: "learning a discrimination entails more than just perception; cognitive factors are also involved. An animal might perceptually distinguish a triangle and circle from the start, but nevertheless requires training to learn that response to one stimulus is followed by reward whereas response to the other stimulus is not. A human subject might require several trials before realizing that a triangle is always rewarded and a circle is not. *But no one would argue from this fact that on these first few trials the subject did not perceive the forms veridically*" (p. 369) [Dretske's emphasis].
18. I put the word "recognition" in scare quotes because this is *not* a genuine cognitive achievement. No *beliefs* are produced by this simple mechanical system—nothing having the intentional struc-

ture of *knowledge*. For more about what constitutes the distinguishing features of a belief state, see section 3.

19. The word “perception” is often reserved for those sensory transactions in which there is some cognitive uptake (identification, recognition, etc.). The sense of the term I allude to here is the sense in which we can see, hear, and smell objects or events (be aware or conscious *of* them) without necessarily categorizing them in any way. This point is more fully discussed below (next note and the following section of this chapter [latter note included herein]).
20. In *Seeing and Knowing* I argued that seeing *s* (a dog, a tree, a person) was essentially nonepistemic: no *beliefs* were essential to the seeing. Although we (adults) typically acquire a variety of beliefs about the things we see, seeing a dog, a tree, or a person is itself a relationship that is independent of such beliefs—one *can* see *s* without believing that it is *F* (for any value of *F*). My present way of expressing this point is different, but the point remains the same. The only modification consists in the requirement that in order to qualify as a perceptual state (see *s*) a structure must be *coupled* to a cognitive mechanism capable of exploiting the information held in the sensory representation. In this respect my present view is somewhat closer to Frank Sibley’s (1971). I am indebted to David Ring for helpful discussion and clarification of this point.

This excerpt from

Philosophical Psychopathology.
George Graham and G. Lynn Stephens, editors.
© 1995 The MIT Press.

is provided in screen-viewable form for personal use only by members of MIT CogNet.

Unauthorized use or dissemination of this information is expressly forbidden.

If you have any questions about this material, please contact cognetadmin@cognet.mit.edu.