Research in Psychophysics

L. D. Braida, Tom N. Cornsweet,
N. I. Durlach, David M. Green, Herschel
Leibowitz, Alvin Liberman, R. Duncan Luce,
Richard Pew, and Carl Sherrick

INTRODUCTION

Psychophysics as a discipline arose during the middle of the last century from the burgeoning success of physics in explaining natural phenomena. It was thought that the methodology used so effectively in dealing with the relationships of material things to each other--at that time being translated into the industrial revolution--could also be used to characterize the relationships between the material and the mental worlds. At the beginning the movement attempted to encompass a wide range of phenomena, from the just detectable difference in weight of objects to the beauty of works of art. The lasting heritage from that era of psychophysics is the acceptance of the scientific procedures of the physical sciences in psychology: fully defined and replicable experimental methods, results expressed in numerical terms, the formulation of theories in the form of mathematical functions, and the use of these theories to extrapolate (i.e., make predictions) to specified new situations.

As it is seen today, the goal of psychophysics is to understand at the most basic level the ways in which one apprehends changes in the physical environment, especially when one is attending to such changes. Often the simplest possible changes in stimulation are used--a shift in the energy level of a spot of light or a change in the frequency of a pure tone. Such simple stimuli are usually referred to as signals. In other studies more complex stimuli are used--the presentation of alternating bands of dark and light on the whole visual field or spoken words. When the stimuli are very simple, we usually speak of visual, auditory, or tactile sensation. As the stimuli become more complex and the questions become more ones of grouping and classifying stimuli rather more than just detecting them or telling them apart, we begin to speak of visual, auditory, or tactile perception. Psychophysics incorporates all of the work on sensation and some of the simpler aspects of perception; however, the boundary line between sensation and perception is rather fuzzy. Reading is a perceptual process that very few would class as psychophysical, but just how one perceives individual letters and even how one groups them into words many construe as part of psychophysics. In sum, then, psychophysics is the study of how one makes distinctions about the energy impinging on one, and it attempts to understand one's abilities to extract information from the environment. When the focus of interest becomes that of extracting the meaning embodied in the stimuli, one is outside the bounds of psychophysics and into the general domain of cognitive psychology.

Is psychophysics just a single name for the study of the several senses? Not really, because there is considerable interest in comparing systems and in finding common mechanisms. Such mechanisms may arise for two reasons. First, nature often repeats its solutions, e.g., the same technique in the peripheral nervous system for enhancing differences appears in several of the senses. Second, much of what is interesting about the way information is extracted appears to be mediated not only by the sensory transducer--ear, eye, taste buds, etc.--but also by the brain itself, and some of the higher processing may use the same brain mechanisms for different senses.

Our discussion is divided into two major parts. In the first we treat basic research in psychophysics and in the second some of its applications. The basic research is separated into four parts: the psychophysics and physiology of the visual system, with major emphasis on the sensory aspect; the same for the auditory system; the common problems that arise when cognitive factors, such as instructions, play a crucial role in how the subject responds to the questions posed; and the study of the motor responses of a person in dynamic interaction with the sensory environment, as in flying. Several of the sensory processes--touch, taste, pain, and heat--are not dealt with explicitly in this first part because considerably less is known about them than about vision and audition. (Some aspects of that work, especially for touch, are described in the section on applications.) The section on applications is concerned with attempts to classify and to overcome serious sensory deficits, except for the first subsection, which is concerned with the problem of why drivers think they can see obstacles better at night than they really can.

Much of the work we are mentioning is inherently rather technical. We attempt to keep jargon to a minimum, but some technical concepts are needed. We often attempt to suggest them by example or illustration rather than offer explicit definitions. For truly precise definitions and statements of results, the reader is urged to consult the literature. One general point of caution. It is essential to maintain a distinction between the physical attributes that are manipulated experimentally and the subjective sensations that are related to these physical variables. For example, as experimenters we can manipulate two features of a pure tone, its amplitude (or intensity) and its frequency. The listener can speak of and respond to questions about its loudness and pitch. Loudness is a sensation that is mainly affected by the signal amplitude, but it is also influenced by frequency, which is why one has a loudness control as well as a volume control on good amplifiers. And pitch is affected mainly by the signal frequency, but it is also affected to a lesser degree by the amplitude. Thus, loudness and intensity are not interchangeable, nor are pitch and frequency. The same is true of luminance and brightness, sugar concentration and sweetness, etc.

BASIC RESEARCH IN PSYCHOPHYSICS

Vision

The optical surfaces of the two eyes interact with light reflected from surrounding objects to form images of those objects on a layer of tissue inside each eye that converts light into electrochemical signals. The images have the same properties as those formed on the film of an ordinary camera, and the chemicals in the eye respond to the light in ways that are closely analogous to the ways in which the corresponding chemicals respond in film, producing what is essentially a chemical picture of the scene. However, within a few thousandths of a second after this picture is formed, the neural elements in the eye begin to transform it through processes that are very different from anything in a camera. Parts of the neural processes are magnified and others are drastically reduced. Separations are made, too, so that different aspects of the image are conducted to and processed by different sets of neural mechanisms. Then the neural images are conducted out of the eye and into various regions of the brain, where further, more complex processing occurs. These neural processes strongly affect what we see in ways that are only now becoming understood.

For example, the fact that a lemon looks yellow and bright whether it is seen inside a kitchen or outside on a tree in broad daylight may not seem surprising. However, if it were not for some very complex processing of signals in the visual system, the color and brightness of the lemon, and in fact of all objects, would seem to change drastically when the intensity of the light falling on them changed. Without such processing, the lemon that looks yellow and bright on the tree might look dark and violet in the kitchen. It is only within the past 20 years that vision scientists have begun to develop a clear understanding of the nature of these processes and of their consequences for human visual perception.

For a while it looked as if the greater insights into the mechanisms of human sensory responses would come from the anatomists and physiologists. Once the technological tools (microscopes, oscilloscopes, amplifiers, computers, chemical probes, etc.) were made available to the "wet" scientists and once they had gotten thoroughly imbued with a mechanistic viewpoint, they proceeded to lay bare the immediate physical substrates to human sensory responses: the mechanical transducers (e.g., the optics of the eye, the middle and inner ear), the sense cells and the way the physical stimulus (light, sound, pressure) activates the cells, the neural pathways and their connection in the brain, etc. Spectacular successes were achieved: It was demonstrated how all of vision is funneled through a chemical stage inside the receptor's cells and that the wavelength dependency of vision could be accounted for by the absorption properties of the chemical visual pigment molecules at different wavelengths. More recently it has been shown that the individual nerve cells in the visual pathways behave in a manner more complicated than merely being photo cells. They respond best to rather specific light shapes, for example, bull's-eyes or bars of light. The excitement generated by these findings was not about the discovery that individual cells in the sensory nervous system behave in a complicated way, but that the kind of complication they exhibit matches the complication that behavioral researchers find in human and animal responses to visual stimuli.

Over the years psychophysicists had observed many kinds of visual phenomena that did not fit in with the view that the visual system operates as a simple set of light transducers. Blurred edges are seen as sharper than they actually are, a feature essential to the successful reproduction of photographs in newspapers or moving images on TV screens. The contrast of adjacent areas seems enhanced. Color values depend on brightness and color context: A piece of gray pottery appears dark on a white tablecloth but light on a dark mahogany table. Spatial patterns appear larger or smaller or even tilted in ways that are predicated on neighboring patterns, a fact that architects for centuries have known they must take into account. Many of these effects are celebrated as "visual illusions," but their more insightful investigators have always accepted them as manifestations of the transformations of sensory signals in their passage within the human central nervous system. The discovery of neural networks with properties that fit almost exactly some of the specifications set by psychophysicists constituted an important landmark.

Although there was general jubilation among sensory physiologists that they had "explained" behavior, or at least validated observations of behavior in mammals, in fact the situation was precisely the reverse: What might otherwise have appeared as an arbitrary arrangement of excitatory and inhibitory connections in the nervous system suddenly fell into place as meaningful once it was realized that it matched the performance of the whole organism.

This interaction between physiological findings in animals and behavioral observations in humans must be regarded as a high point in the investigation of sensory phenomena. Matches were sought between the ability of an observer to distinguish colors, shapes, movement, and three-dimensional depth of targets, on one hand, and the inbuilt selectivity for such features of cells in the retina and sensory brain, on the other. Cells that tend to respond predominantly to a line of a certain tilt, for example, are said to have this line as their "trigger feature," and the dissection of animal behavior into channels delineated by trigger features of cell classes in the incoming stream of information seemed to be the obvious approach. All that was needed was to enumerate fully the cells' trigger features and describe their properties in detail, and, in principle, the sensory process would be characterized.

As this process was pursued during the last decade, a difficulty emerged. It concerned the first step in the process, the enumeration of cell classes by their trigger features. Single-cell experiments in mammals are difficult and the search for the precise trigger feature of a cell is time-consuming. In fact, the possible combinations of visual stimulus parameters in the domains of space, time, brightness, and chromaticity is astronomical, making it impossible to proceed in a systematic manner through all possible target positions, shapes, velocities, colors, etc., in order to state the cell's adequate stimulus with assurance. Where anatomy helped, as in the retina or the first relay in the brain, the approach was moderately successful, but now that the search is being continued into secondary and tertiary cortical projections, results are becoming more and more ambiguous.

Thus there is a need for organizing principles to be brought into the single-cell laboratory or, if you will, preconceived ideas as to the likely combination of stimulus parameters that may constitute the trigger features of a cell. And here psychophysics reemerges as a guide. As the routing of sensory signals within the central nervous system becomes fuzzy, with many different loci of activity, the pattern may be clarified by knowing what the major response modes are of the whole organism.

Over the years steady progress has been made in the delineation by psychophysical means of the modes of signal processing by the human nervous system. A great deal of the most reliable data is obtained by threshold experiments, i.e., by the determination of the least stimulation that can be detected by the human nervous system. A threshold, as the word suggests, is a boundary or dividing point. When we speak of a light threshold we mean the energy level of the light such that above that level it is seen or sensed; below that level it is not. Threshold experiments are ones that measure this energy level for a wide variety of stimuli. By imaginative manipulation of stimulus variables, enormous areas have been charted. For example, the minimum quantity of light needed in order to be detected, called the visual threshold, is now no longer being measured just for a spot of light: The spot can be placed on backgrounds of various sizes, shapes, and colors, and in this way the confluence of excitation from neighboring regions (in time, space, and color) can be ascertained. Thus it is possible to predict whether certain light signals are visible in a whole range of road, rail, and aviation traffic situations. Instead of using spots of light, thresholds can be determined for more complex and interesting shapes--and it has been found that detection can be improved by the judicious selection of the overall shape or size of a pattern. The same applies to the velocity of targets.

A particularly interesting approach is to adapt the visual system to a particular stimulus situation by having a subject view it for perhaps a minute. When, for example, a band of stripes of given spacing is viewed for a while, other stripe spacings look wider or narrower than before, and they require a greater luminance ratio between the light and dark stripes--a higher contrast--to be detected. In finding the range of stripe widths affected by adaptation to a particular pattern, it becomes possible to describe the channeling of information in the visual sensory system. Stripes or other patterns whose visibility is unaffected are presumably not processed in the same channel; conversely, the extent to which any pattern is affected helps to outline the characteristics of the channel.

Channels--three in number--to transmit color information have been known to exist for a long time, but the description of channels used in spatial vision and in the velocity domain is relatively recent. Channels can be quite specific. For example, it has been shown that adaptation to the back-and-forth jitter of a pattern leaves unaffected the detection of zooming motion--signifying the existence of several separate motion channels. (For a detailed application of these results, see below.)

Another area in which modern psychophysics is pointing the way is in outlining the ultimate sensitivity of the organism to small changes in stimulus situations. Time discrimination in hearing can be in terms of microseconds, in vision a fraction of a millisecond. The eye can resolve scenes as well as can be done by any diffraction-limited optical instrument of the same aperture. Shades of color can be discriminated by most observers in a way that taxes the capability of colorimetric schemes. In terms of the sensitivity of the eye to light, one or two light quanta arriving simultaneously are enough to be seen; no photo cell can do better. The localization ability of the human visual apparatus is so fine--a few seconds of arc or one hundredth of a milliradian--that it is called hyperacuity to set it off from ordinary telescope resolution. All these fine discrimination capabilities of the human visual sensory apparatus betray the operation of neural circuits that transcend by a factor of 100 the current best estimates of what nerve cells by themselves can do. They therefore point to the next task of neurophysiology: the elucidation of computational machinery in the brain that accepts relatively crude sensory input and processes it, by means as yet unknown, to yield signals of exquisite refinement. How they are stored (memory) and compared with each other (cognition) is the subject of other facets of behavioral science.

Audition

Psychoacoustics, the psychophysics of auditory phenomena, owes much to two pioneers who initiated the earliest investigations in this area. Lord Rayleigh, the English scientist, almost single-handedly developed the modern field of physical acoustics and also began the scientific exploration of binaural hearing. H. von Helmholtz, the German physicist and physiologist, pioneered the first systematic analysis of auditory phenomena. Their discoveries, theories, and speculations resulted in the field of psychoacoustics, which relates what one hears to the physical properties of the sound stimulus.

An important aspect of any sound is its intensity, since as intensity is varied the loudness of a sound changes. Rayleigh in 1882 invented a disc, now named after him, which provided the first means of measuring physically the intensity of a sound field. Prior to that time the frequency of a sound was the physical parameter of major interest, because changing the frequency of a source altered the pitch of the sound, and frequency could be measured physically with some accuracy.

One of the earliest questions involved how the sensation of pitch is conveyed within the nervous system: How are the major physical aspects of sound, frequency and intensity, coded? In 1863 Helmholtz proposed what became known as the resonance or place theory of hearing. He suggested that somehow the different places along the cochlear duct of the inner ear, where the sense organs of audition reside, are differentially responsive to different frequencies. The analogy is a harp. Just as different strings of a harp vibrate "sympathetically" to different frequencies, so different places along the cochlea vibrate to different frequencies of sound. Place thereby codes the frequency of the sound, and the amplitude or vigor of responses (e.g., how many times the nerve fires per second) codes the intensity of the sound vibration at that frequency.

This theory was a natural extension to audition of the classic Young-Helmholtz theory of color vision. The eye contains three different receptor types, which are differentially responsive to long, medium, and short wavelengths. Thus the color quality is thought to be coded by which receptors are active; the vigor of their total activity codes intensity. The major difference between the place theory of hearing and the Young-Helmholtz theory of color is that, instead of three receptors, there are hundreds of different qualities (different neural fibers) representing the pitch of the sound.

The assumptions of place theory were comfortable to the physics community in part because they nicely meshed with the remarkable mathematical insights of J. Fourier. Fourier's theorem asserts that a complex periodic wave can be represented as a sum of simple sinusoidal vibrations. Place theory assumes that the ear performs such a decomposition, each place resonating to a distinct sinusoid, and the collection of places thereby representing the complex periodic sound.

Early psychoacoustics was not fully systematic because precise means of generating sound stimuli were not available. Typical sounds were tuning forks or crude sirens. Modern psychoacoustical investigation began when reliable instruments for generating and controlling the physical stimulus became available. This era began with investigators adapting the new electronic technology to the generation of sounds via headphones (invented by Alexander G. Bell in 1876) and progressed with later improvements in headphones and loudspeakers. The earliest systematic studies were carried out, not surprisingly, by the Bell Telephone Laboratories about 1930. Their interest in psychoacoustics arose because the ultimate arbiter of the quality of any acoustic transmission system is the sense of hearing. About the same time G. von Békésy, a Hungarian telephone engineer, began studying hearing for the same reason. His investigations later earned him the Nobel prize.

The earliest psychoacoustic experiments of the modern era (e.g., those of R. L. Wegel and C. E. Lane in 1928) studied the way in which a tone of one frequency could make a tone of another frequency difficult or impossible to hear. This effect, called masking, is important since, to the degree that masking is effective, communication is impossible. Information on masking has practical applications in understanding how we hear, or fail to hear, in noisy environments. Information about masking contributed in important ways to the development of effective radio communication systems in aircarft in World War II.

Gradually the model of the ear as a series of resonant channels, each responsive to a slightly different frequency, became a widespread idea and was used to explain a variety of masking results. This general notion was completely consistent with Helmholtz's earlier resonance theory. A most important related phenomenon is called the critical band of frequencies. The basic idea is that when two frequencies are sufficiently close together that they both activate the same resonant channel, then they interact in a way that is quite different from the case of two more widely separated frequencies that activate different channels. For example, consider what happens to the detectability of a pure tone in a narrow band of noise centered around the tone. At first, as one increases the width of the noise band, the detectability of the signal decreases because the total disturbance in that channel grows with the number of components present. But once the band is so wide that any increase affects different resonant channels, then the detectability does not change any further. This is one way to estimate the width of the band. Because many interesting psychoacoustic phenomena are understandable in terms of critical bands, their exact nature has been the focus of considerable experimental and theoretical work.

About the same time as the ealy psychoacoustic experiments, Békésy began to explore the anatomy of the cochlea and was eventually successful in seeing the vibration of the basilar membrane, the delicate tissue on which rest the hair cells, which are the receptor elements of the auditory sense. The amplitude of the vibration is very small; even at enormous intensities the amount of movement is barely detectable using the option methods Békésy employed. In the 1970s much more subtle techniques have been used to study these vibrations (e.g., the Mossbauer technique and laser interferometry). By and large the latter measurements completely support Békésy's earlier observations, confirming that each place along the membrane responds only to a certain narrow range of frequencies, just as place theory would have it.

Following this earlier work, a large number of psychoacoustic investigations explored the finest discrimination that could be made of a small change in a basic physical parameter of the stimulus. One could, for example, hear a change of less than 10 percent in intensity, a change of less than 3 percent in frequency, and a change in angle of spatial locus of less than 1 percent if the source was located straight ahead and emitting a broad spectrum. All of these studies of auditory acuity were aimed at inferring something about auditory processing from a measurement of the smallest detectable change. As in the visual system, many potential hypotheses about auditory processing are untenable because they would not result in sufficient sensitivity to accord with the measured sensitivity of human observers. Many and varied experiments of discrimination capacity continue to be pursued. These new experiments refine and limit the number of reasonable hypotheses concerning the details of acoustic processing.

Animals other than humans have also been studied because in some cases their sensitivities exceed those of humans. As their response to special ultrasonic whistles demonstrates, dogs and cats can hear higher frequencies than can people. More remarkable still, bats navigate and catch prey using the reflections from the ultra-high-frequency pulses that they produce. Bats used the principles underlying sonar and radar millions of years before humans discovered them.

Meanwhile, physiologists continued their exploration of the hearing mechanism in an effort to understand the details of auditory processing. A major breakthrough was the ability to record, electrically, from a single auditory fiber in the acoustic nerve, the VIII cranial nerve bundle. These recordings revealed that each fiber is maximally sensitive to only a narrow band of frequencies, the width of the band increasing with the intensity of the tone. In effect the fiber acted very much as a resonance filter, maximally sensitive at one frequency but capable of responding to other frequencies if the intensity is sufficiently large. In a plot of intensity versus frequency, the curve separating the region of responsiveness from that of no response is called a "tuning curve," and the measurement of a fiber's "tuning curve" is now an essential first step in any serious study of the auditory nervous system.

The tuning curve tells us which place along the basilar membrane contains the hair cell that drives this fiber. Tuning curve analysis is, in short, completely consistent with Helmholtz's place theory. In recent years psychophysicists have devised a masking experiment that generates a resonance-like curve closely resembling the physiologically measured tuning curve. In addition, near the edges of the tuning curve one finds frequency and intensity combinations that appear to produce suppression, a process that can cancel the effects of masking. Again, both physiological and psychophysical data show strong similarities. The observation of suppression suggests a region of inhibitory action flanking an excitatory center, reminiscent of a lateral inhibition mechanism such as that found in vision and the skin senses (see the section on visual modulation transfer function and visual disorders).

Although the place-resonance theory of Helmholtz is supported by all of the available peripheral data, it is not the whole truth. The perception of pitch, especially the pitch produced by complex periodic stimuli such as those arising from musical instruments, is not as simple as the original theory would have it. About 1940 J. F. Schouten of the Netherlands reported some experiments in which subjects heard a low-frequency pitch, called the residue pitch, as a result of certain combinations of high-frequency components with no energy whatsoever at or near the perceived pitch. Thus, although activity at one place signals the corresponding pitch, that same pitch can be produced by the combination of activities at other places. Moreover, as several experiments demonstrated, the low pitch is not the result of nonlinear distortion. This is an auditory illusion in the same sense that there are visual illusions. One is hearing something quite different from what is in fact present in the stimulus. And as with visual illusions, it is important for two different reasons. First, it places a very strong constraint on theoretical ideas about how the auditory system works. Second, it is an illusion that can affect the practice of engineering acoustics. For example, noise engineers were initially baffled by the fact that people complained of noise at low frequencies in jet engines when there was little or no energy at these frequencies. The fan blades were creating very regular patterns of energies at higher frequencies that created significant residue pitches. A decade of further work, both in this country and in the Netherlands, has further clarified and confirmed Schouten's original observations concerning residue pitch. The current consensus is that residue pitch undoubtedly represents the action of some more central process interpreting and integrating the peripheral sensory information. The recognition and acceptance of the facts of residue pitch have also been important in advancing understanding of nonlinear phenomena in hearing. Such nonlinear effects are important and ubiquitous but were largely misunderstood by early investigators.

In summary, exploration of the auditory sense has been one of mutual support among investigators working in physics, physiology, and psychoacoustics.

Cognitive Factors in Psychophysics

Despite the highly successful interplay of psychophysical and neurophysiological research, there are major psychophysical phenomena, some of which were only fully recognized in the past 30 years, that have completely eluded physiological clarification. These have to do with phenomena of the central nervous system--sometimes called information or cognitive processing--that are only partially influenced by the peripheral information arising from the physical signals. Among the topics in this area are: the ability of a subject to attend differentially to aspects of the stimuli impinging on him or her; the trade-off that exists between failing to detect a signal and falsely responding that it is present when in fact it is not; the trade-off that exists between accuracy of performance and the time it takes to respond; how varying a signal or its surrounding environment affects the subjective growth of sensation (e.g., what does it mean to say that the noise level has been reduced by half?); and the inability of most people to identify correctly more than about seven signals that differ along just one physical dimension.

We elaborate several of these examples. Consider driving on a lonely road. One is continually scanning for danger signals--another car, fixed obstacles such as trees or rocks, and of course pedestrians. Having selected a speed at which to drive, there remains another variable under one's control--how reactive to be to apparant signs of danger. If one is very reactive, applying the brakes at the first partial indication of danger, then frequent false alarms result (i.e., braking when there is no obstacle) but less danger of striking something. If one is less reactive (i.e., waits until a clear danger is present before applying the brakes), then fewer false alarms may result but the chance of an accident rises. Clearly, the driver has available some freedom to decide just how reactive to be, both in terms of the amount of evidence collected in a fixed time that is sufficient to cause braking and in terms of the amount of time to delay before making the response. Such trade-offs are ubiquitous in sensory psychology. Considerable work has gone into their study, and rather elaborate mathematical models have been developed in an attempt to capture the basic principles of the decision processes that are involved.

This work has shown that questions such as: "How fast can a person respond to a signal of such-and-such a character?" or "How likely is it that such a signal will be detected in a certain environment?" are either meaningless as formulated or, at the very least, require very subtle answers. How fast one responds depends greatly on how many false and/or anticipatory responses are permitted; how likely one is to detect a signal depends on how many false alarms are tolerable. With exactly the same signal conditions and the same subject, the likelihood of detecting a signal can be varied from nil to certainty simply by varying how likely a signal is to occur in a fixed time period or the nature of the rewards for correct responses and the punishments for the two types of errors. The past several decades have provided us with much data and sophisticated models concerning the nature of this trade-off.

Loudness grows with the amplitude of the sound wave, visual size with physical extent, and the sensation of shock with voltage. But exactly how do these sensations grow? This question was initially raised during the 19th century; G. Fechner's attempt to answer it led to the beginnings of psychophysics, and it has been inextricably intertwined with all subsequent theoretical developments in the field. One of the more striking of these emerged in the 1960s from results obtained by S. S. Stevens, who simply had subjects assign numbers to signals in proportion to the subjective sensations they engendered. It turns out that they can do this seemingly impossible task with considerable regularity. As a first approximation, sensation measured this way grows as a power of signal amplitude, A^{β}, where A denotes amplitude and β is an empirical exponent. The exponent involved varies from less that the cube root (loudness and brightness), through the linear function (line length), to something in the neighborhood of the cube (electrical shock). There is much work going on attempting to understand the theoretical basis of these relations, to understand how various factors affect these subjective sensations, and to understand how they combine.

Some specific questions are these. Suppose that a sound is composed of several pure tones. Can the loudness of the combination be predicted from the separate loudness of the components? From what was said in the section above on audition, it comes as no surprise that the answer is much affected by the structure of the critical bands.

Given separate sounds to the two ears, how does the overall sensation of loudness depend on the individual loudness in each ear? Consider the apparent size of the moon. We are all familiar with the fact that the moon rising over the horizon seems huge compared with its size when it is high in the sky; subjective estimates yield a factor of about two. Yet from photographs or by viewing the moon through a mailing tube we know that the visual angle at the eye is virtually identical in the two cases. So, why does it seem larger at the horizon? No one really knows, despite the fact that the phenomenon has long been recognized and a number of attempts have been made to explain it. Such an illusion, one of many of which we are aware, is surely not an isolated curiosity; rather it tells us something about the basic information processing carried out by the brain on the input signal.

Another phenomenon that has been recognized since the 1950s but whose basis is still not fully understood is this: If one is asked to identify which of two sounds, one twice as intense as the other, has been presented, one can do so with perfect accuracy; sounds separated by that amount are never confused. Suppose the total number of sounds is increased to 10 and successive ones are still spaced by the same factor of 2; errors will then be made, e.g., sound 8 will sometimes be called sound 7 or 9, despite the fact that all can be perfectly discriminated as pairs. The data tell us that the ability to identify sound 8 and to discriminate it from 9 depends on other sounds that might be presented. The same is true of brightness and most other modalities. (The major exception is auditory frequency: Some people exhibit the phenomenon of almost perfect pitch.) As far as anyone knows, the brain has exactly the same peripheral information when a particular signal is presented, whether it is 1 of 2 or 1 of 10, yet once the number exceeds about 7 there is increasing difficulty and confusion in identifying which signal is presented. Why? Various theories have been offered, but to date none seems fully adequate. It is clear that as yet we do not fully understand the nature of the coding involved and the processing done by the brain.

We mention these easily demonstrable conundrums because they are both very familiar and very difficult to understand. There can be little doubt that the brain, when processing even the simplest of signals, is using information that extends well beyond the particular signal at hand. Psychophysicists are making attempts--often mathematical ones--to formulate the types of processes that may account for these phenomena.

Skilled Performance

Another important area of basic research is the development of quantitative models for how a person uses visual or auditory information in combination with his or her cognitive skills to control and manipulate machines.

When one drives a car or flies an airplane, the senses, particularly the eyes, take in information from the environment, and the brain processes the information, together with goals or intentions, to send signals to activate the muscles. The muscles in turn move the steering wheel or control stick in order to control vehicle movement. When a person serves as a vital link in a control system, it becomes very important to be able to understand the human behavior involved in engineering terms. The sensory-motor response of the driver or pilot affects the system's overall stability and performance, just as do the tires or ailerons.

To the naive observer the task of driving seems to be automatic and to take very little mental effort. In fact, much of the activity involved appears to be unreportable. The individual performing the control task cannot describe how he or she does it. Research has shown, however, that it does require mental effort and, in fact, in most situations such control employs processes that would be described as distinctly cognitive.

A pilot following the glide slope needle that directs the plane toward the runway is performing the simplest of tracking tasks. However, the choice of a control action entails selecting a sequence of muscle commands that must take account not only of the stiffness and inertia of the pilot's own arm but also the sluggish dynamics of the plane itself. The response of the plane is so subtle that one cannot depend on visual cues or pressure on the seat of one's pants to provide adequate feedback. One must integrate these cues with a learned prediction of how the vehicle will respond. We say the pilot has an internal model of the vehicle dynamics and incorporates its response into his muscle command planning.

This point is made forcefully when a tire blows out or a yaw damper fails and the handling characteristics of the vehicle change suddenly. The control behavior appropriate to the system prior to the emergency is quite different from that required after the change. The internal model must be changed and changed rapidly.

When a pilot executes a preplanned maneuver such as a turn to final approach, the cognitive demands are even greater. Since no pathway in the sky exists, the basis for executing a particular turn must be drawn from memory. Some argue that a particular maneuver comes from a schema representing turns in general made particular to the set of conditions found when the turn is begun. The execution also must take account of the specific vehicle characteristics and the nature of the controls involved.

The first attempt to represent human control behavior in such engineering terms was accomplished in 1947 by A. Tustin, a British scientist. The first extensive set of data that described tracking behavior in terms of control engineering equations was completed in 1956 by J. I. Elkind. That work represented a significant advance in experimental measurement of dynamical systems as well as a landmark contribution to the literature in engineering psychology.

Since that time modern control theory has extended the power and complexity of the systems that can be analyzed. The optimal model of manual control describes the behavior in terms of state estimation, information processing, and response generation. It is interesting that the model also assumes that the driver or pilot employs an idealized internal representation of the system being controlled, exactly analogous to the cognitive models described in an earlier section.

Research on manual control and the prediction of human motor performance still has difficulty predicting the effects of learning on performance. There are special difficulties in predicting those practice effects associated with voluntary maneuvers for which there is no explicit pattern to be followed. Manual control modeling research is also being broadened to include prediction of the behavior of multiperson crews when the task of actually controlling the vehicle is overlaid with a myriad of decision-making and procedural tasks, sometimes referred to as supervisory control. One such model describes the activity of the three-person crew of a commercial jet transport during approach and landing. Others have been designed for bicycle riding and motorcycle handling.

Besides the conceptual understanding they contribute, manual control models have been practically useful, for example, in predicting critical design conditions that need careful experimental study. The models of jet transport crews were developed to evaluate alternative landing procedures and staffing requirements.

As this discussion reveals, research on manual control has provided a most interesting and stimulating interchange between experimental psychologists and engineers. It has captured the creative talents of researchers from several fields, not only in terms of the development of new theoretical concepts and measurement procedures, but also in translating these concepts into data and methods useful in the system development process. When analytic methods are available to predict performance, it is possible to narrow the range of candidate designs that need to be evaluated in depth. While it would be impossible to estimate the cost savings attributable to these developments, it is clear that they have substantially reduced the development time and the costs associated with experimental evaluation of alternative designs.

EXAMPLES OF THE INTERPLAY BETWEEN BASIC RESEARCH AND APPLICATIONS

In discussions of the relative importance of basic and applied research programs, an economic analogy is often made. Basic research is the savings account, to be accumulated and husbanded against the lean years, when the checking account of applied research runs low. The cash reserve, it is said, is drawn upon to replenish the drained resources of the experts in application.

The analogy is incomplete in this sense: It neglects the vital presence of those whose activities move the masses of information in and out of the fund of knowledge. Basic knowledge can no more be static than the funds in a savings bank can if there is to be any gain from it. But a requisite for moving, altering, and adding to information is people who are proficient in handling and remodeling it. What William James spoke of as the cash value of an idea is tangible only in the hands of a skilled barterer.

We turn now to this interplay of basic and applied research. Several examples of applications, one inexorably intertwined with the basic research, have already been mentioned. In what follows we take up a number of additional cases in which basic psychophysical research led or is leading to significant applied work. In some cases, the application is complete. In others, the basic work has suggested an idea for solving a problem and the work is currently being pursued. As always, promising routes seem so until they fail; few roads end at the goal, and most are a lot rougher and not nearly so straight as anticipated.

Two Modes of Visual Processing and Night Automobile Accidents

For some time psychologists studying visual perception have posited two independent and dissociable modes of processing visual information. The focal mode is concerned with object discrimination and identification or, more generally, the question of what. It is subserved primarily by the cortex and is typically well represented in consciousness. Because focal functions involve the higher spatial frequencies, i.e., finer visual textures, they are optimal in the central visual field and are systematically related to both luminance and refractive error. The other mode of processing, referred to as ambient vision, is concerned with spatial orientation or, more generally, the question of where. The properties of the ambient and focal modes differ along many dimensions. Although spatial orientation is certainly possible, if not superior, with the central visual field, it is adequate with stimulation of the peripheral retina in spite of the coarse resolution properties of the latter. Coarse patterns are sufficient for ambient functions, so they are less sensitive to both refractive error and luminance than are focal functions. With respect to consciousness, the ambient system is often poorly represented, although by directing attention one can be aware of ambient activity.

A number of recent ablation studies, as well as observations of brain-damaged people, have suggested that it is possible for some orientation ability to be spared despite loss of focal vision. L. Weiskrantz has referred to this interesting phenomenon as "blindsight." For our purpose the fact that it is possible to walk while reading demonstrates the dissociability and some basic characteristics of focal and ambient functions. Even though attention is dominated by the reading material, orientation in space is carried out confidently and accurately by the peripheral visual fields operating at an unconscious or subconscious level. If illumination is lowered or the retinal image blurred, the ability to read is degraded but orientation is relatively umimpaired.

A critical problem in vehicle guidance, which can be understood in terms of the theoretical approach, is the high frequency of nighttime driving accidents. Automobile accidents, of course, have multiple causes. The role of illumination is demonstrated, however, by studies indicating that, when other factors are held constant, accidents, particulary those involving collisions with pedestrians, increase dramatically under lowered illumination. It is well known that under twilight and nighttime conditions many visual capacities, such as spatial resolution, stereoscopic depth perception, contrast discrimination, and reaction time, are degraded. This is reflected in analyses of nighttime accidents in which drivers frequently report that they did not see a pedestrian or other obstacle in time to stop. In some cases, the sound of impact was heard before the driver was aware of the pedestrian. What is curious is that drivers typically do not reduce their speeds at night, even though they are probably aware through personal experience, or even through knowledge of literature, that their vision has been degraded.

A possible explanation for this paradox may be derived from considering the two modes of processing. Driving an automobile, like walking, flying, and sailing, involves two parallel tasks. Spatial orientation is accomplished by steering the vehicle, which requires continuous evaluation of the location of the vehicle relative to the road. In terms of the two modes of processing, steering is concerned with where and is an ambient function. Driving also involves focal vision, the role of which is to monitor the roadway ahead for pedestrians, other vehicles, and obstacles, to read traffic signs and monitor signal lights, and to judge the distance and speed of other vehicles. In daylight both the ambient and focal modes are operating at their maximal capacities. Under twilight and nighttime conditions, however, there is a selective degradation of the two modes. Focal visual functions are degraded, i.e., spatial and stereoscopic acuity are reduced and contrast sensitivity is diminished. (For many individuals the ability to appreciate detail is further degraded by a condition known as night myopia.) The efficiency of ambient visual functions, however, is not reduced by lowered illumination. As long as minimal visual stimulation is available, it is possible to steer the vehicle adequately. In terms of the performance information available to the driver, it is the ambient mode that dominates. Since the demands on focal vision are intermittent, information about the degradation of focal vision is only rarely reflected in the operator's performance. As a result, the driver is not aware that there is a problem with the degradation of focal vision and therefore typically maintains the same velocities at night as during the day.

As is often the case, understanding the basic cause of a problem suggests methods for amelioration. To reduce the high nighttime accident rate a number of possibilities are apparent. Obviously, illuminating highways would be expected to be effective, and this is supported by empirical observations. However, economic considerations limit this possibility. Other alternatives are to post different maximum velocities for nighttime and daytime driving conditions. Before the introduction of the uniform national 55 mph speed limit in the United States, only a few states followed this practice, usually on major highways. To our knowledge, different speed limits have not been posted in areas where degradation of focal vision would be expected to play a role in accidents involving pedestrians. Another possible measure is to educate drivers regarding the potential dangers associated with the selective degradation of vision at night. This procedure would be expected to be particularly effective for younger drivers, whose habits have not been established. The implications of selective degradation should also be communicated to pedestrians and cyclists, who should be encouraged to take special measures to increase their visibility at night in order to compensate for the loss of focal vision of drivers.

Visual Modulation Transfer Functions and Visual Disorders*

A powerful tool for studying visual processing is the measurement of what is called visual describing function or, more loosely, visual modulation transfer function (MTF). Briefly, the procedure for studying the visual MTF is like this: A person looks at a television screen on which a pattern of stripes is present. The experimenter determines how much contrast the stripes must have in order for the person to be able to detect them. This just-visible contrast depends strongly on the spacing of the stripes, on whether they are flickering and if so at what frequency, and also on the colors of the stripes and the condition of the subject's eye.

In general, measures of the MTF provide significant amounts of information about the general behavior of the human visual system and, together with data from physiological experiments on the eyes and brains of various animals, are rapidly leading us to a good understanding of the anatomical and physiological structure of the human visual system. To give just one example, it is now clear that certain neural structures in the human retina are organized in such a way that when a small spot of light falls on the retinal surface, it not only produces an increase in the activity of a few nerve fibers leading from the eye to the brain, but it also produces a decrease in the activity or responsiveness of all of the nerve fibers corresponding to adjacent spots on the retina. This is called lateral inhibition. We have good estimates of how strong these effects are, the distances across the retina over which they operate, their sensitivity to the timing of changing light intensities, how these factors vary at different parts of the retina, and many other similar parameters.

This kind of information is important in a number of ways. First, and perhaps most important, it provides strong hypotheses about how the brain itself works, because the retina is closely related to much of the brain in its embryology and structure. Second, it provides explanations for many visual phenomena, such as the fact that objects look the same regardless of the intensity of light illuminating them, which has puzzled scientists for hundreds of years. Third, it permits accurate predictions of the appearance of unfamiliar patterns of light, such as those experienced by astronauts when traveling in outer space.

The use of the MTF in visual psychophysical research has also been valuable in advancing our understanding of the perception of complex patterns such as letters and faces. Any complex pattern can be broken down into a series of component stripe patterns, each with a specific width, orientation, contrast, and position. By viewing the visual system as a series of filters, each sensitive to a different set of stripe widths and orientations, accurate predictions concerning the perception of complex objects can be made once the component stripe pattern used to construct the complex pattern is identified and the MTF of the visual system is measured. This approach has been

---

*The authors express their thanks to Jane E. Raymond for valuable suggestions for this section.

valuable in basic research on object perception and recognition, and it also has been successfully applied to practical issues.

Although the concepts surrounding the visual system MTF have been useful in areas such as aviation engineering and electronic visual communications engineering, the most widely explored applications of the MTF to date have been in the area of clinical medicine. Three examples follow.

Contrast sensitivity functions were first used clinically by neurologists investigating patients with disorders of the central nervous system such as cerebral lesions, epilepsy, and multiple sclerosis. As a diagnostic tool, measurement of the MTF has been particularly successful in detecting visual involvement of multiple sclerosis in patients who appear normal on ophthalmological examination. Aside from the diagnostic value of this procedure, the data obtained can also be employed to indicate a physiological rather than psychogenic basis for visual complaints and perceptual difficulties.

Second, it has been known for centuries that those children with eyes that do not point in the same direction exhibit characteristics different from those who need glasses but have not worn them. Differences are also found in their MTFs, suggesting differences in the mechanisms of visual loss. Study of the differences may lead to improved methods of prevention.

A third example of the application of MTFs to visual pathology is in the detection of glaucoma. During 1979 and 1980, it was shown that the MTFs for glaucoma patients and many glaucoma suspects differ in characteristic ways from those of people with normal vision. These differences may provide an important means of detecting glaucoma before it has caused serious visual loss. Equally important, the nature of the MTF differences between glaucoma patients and people with normal vision provides important evidence about the actual pathological processes that occur during the course of the disease.

In summary, the study of how the visibility of a pattern of stripes is affected by various characteristics of the stripes and by properties of the eye is leading us to a better understanding of the nature of both normal and pathological human vision.

Speech Perception and Reading Machines

Our understanding of speech perception, as well as our ability to put that understanding to practical use, derives from the confluence of several currents in science and engineering. We trace one such current here, one that exhibits the interplay of a practical problem and the basic research in psychophysics that contributed to its solution. The example is striking, because the practical problem seemed at the outset to require very little more knowledge than was available at the time and also because the basic research, once undertaken, has proved useful beyond the demands of the particular problem that provided the initial stimulus.

The problem, first seriously attacked at the end of World War II, was to build a reading machine for the blind, a device that could scan print and produce an understandable acoustic signal. Some, perhaps all, of those who set out on this enterprise were guided by the assumption that the device had only to produce, for each letter, a pattern of sound that was distinctively different from the pattern of all other letters. Blind users would presumably be able to read after learning to associate the sounds with the letters. The rationale for this device was to be found in an obvious fact and in a seemingly obvious assumption about that fact. The fact is that sound and the ear work well together for the purpose of conveying in speech the very phonemes that the letters of the writing system (approximately) represent. The seemingly obvious assumption was that the sounds of speech are related to those phonemes in a straightforward way, much as the sounds of the reading machine would be related to the letters, and that, in this crucial respect, perception of speech was assumed not different in principle from the perception of any other sounds. Accordingly, it was expected that the sounds of the reading machine would work as well as speech, provided only that they were distinctive and that the users were given sufficient training.

In fact, no arbitrary sound alphabet could be made to work well, no matter how distinctive its individual elements or how long the training of the users. As the conclusion took shape, some of those engaged in the undertaking began to suspect that the assumption about speech and its perception was wrong. There is apparently something special about speech, something that the arbitrary sounds of the reading machine could not capture. Thus it happened that some investigators put aside their work on the reading machine in order to take up the basic research on speech perception that was required if they were ever to find out in what ways speech is special.

At the outset, the task of studying the perception of speech was no different in principle from that of studying the perception of anything. The first step is to find the cues--the physical stimuli--that control the perception; more generally, of course, the aim is to characterize the nature of the relationship between those cues and the precepts they support. Nor is the research procedure different in principle from that which had always been followed in perceptual psychophysics. It was to use methods of analysis to formulate working hypotheses about the cues and then to test those hypotheses by synthesis--that is, by manipulating the physical signal in ways appropriate to the hypotheses so as to determine the effects on the sound as perceived.

But applying the standard procedure to speech was complicated by several special difficulties. For one, many of the most important acoustic cues for speech are in the dynamic aspects of the acoustic pattern, a circumstance that imposes special requirements on the development of an appropriate research synthesizer. A further difficulty lay in the fact, not fully understood when the research began, that the relationship between acoustic signal and phonetic percept is peculiar in ways that make it significantly harder to see just what form the cues may take and where in the signal they may be found. As a result a considerable amount of time had to be spent in designing and perfecting an appropriate research synthesizer. The synthesizer needed to provide control of the putatively relevant aspects of the signal, including especially those of a dynamic character, and also needed to be sufficiently convenient so as to permit the very large amount of experimentation that proved necessary to disclose the special nature of the relationship between signal and phonetic percept.

As the research progressed, many acoustic cues were found and their effects evaluated. With that done, it was possible to see some of the general characteristics of the relationship between the cues and the percepts. Perhaps the most important of these is that the acoustic segmentation does not correspond, as had been originally supposed, in a straightforward way to the segmentation of the phonetic message. Rather, the information about any particular phonetic segment is widely distributed through the signal and overlapped, often completely, with acoustic information appropriate to other segments, thus reducing the number of acoustic segments per unit time that must be perceived. A simple but effective way to demonstrate

this high degree of interdependence is to try to substitute identical phonetic elements from other contexts. For example, suppose one tries to synthesize on a tape recorder the word <u>cat</u> by splicing the <u>c</u> from <u>car</u>, the <u>a</u> from <u>has</u>, and the <u>t</u> from <u>cut</u>. The result is unintelligible gibberish. This is so because, in production, the phonetic segments are coarticulated, with the result that information about several successive segments is normally encoded into and transmitted simultaneously on the same parameter of the acoustic signal. To disentangle those segments in perception requires a specialized process. But if one possesses that specialization, as humans do, then the parallel transmission of segmental information that characterizes the speech code makes it a uniquely effective way of communicating language by sound.

Once the nature of the speech code was understood, it became possible to synthesize speech by explicit rule and thus to have, in principle, an important component of a reading machine that produces, not arbitrary sounds, but speech. To synthesize speech by rule means that, beginning with an input string of letters (or, more properly, their phonetic transcriptions), one produces speech by automatically applying explicit encoding instructions of a kind that can be dealt with by a computer. As first produced in 1958 (though not then by computer), speech synthesized by rule was at best intelligible only in the narrow sense of the word. Listening to it took an effort. Indeed, the extra attention needed to perceive the phonetic aspect of the message was often such as to defeat the attempt to grasp its meaning. But progress continued, and now the speech that can be synthesized by rule is much improved.

A reading machine for the blind is rapidly becoming a reality. Such a machine requires, in addition to rules for synthesis and the means for implementing them, an optical character reader to identify the printed letters and a method for converting English spelling to the phonetic transcription that keys the synthesis. These components exist and can be appropriately linked, so a blind user has a way to convert print to reasonably intelligible speech. One may wonder, of course, whether the devices currently available are good enough for all purposes. Do they, for example, still require too much effort of one who wishes to read large quantities of difficult text? Further research and development will almost certainly produce further improvements.

Automatic synthesis of speech has application not only to the blind but also to some people who are, for any of several reasons, unable to talk. In business, industry, and education, speech synthesis is finding increasing application in human-machine interactions of various kinds. Moreover, the knowledge gained through research on speech perception has proved useful for purposes other than the synthesis of speech. It provides essential information for attempts to build automatic speech recognizers, and, in particular, for understanding the nature of the difficulties that must be overcome if such machines are ever to be as versatile as we would like them to be. It furnishes the key to understanding why it is that beginning readers find it so difficult to develop an explicit awareness of the way that words can be segmented into the abstract units that the alphabetic letters represent. It helps us understand the particular nature of the difficulties that the hard-of-hearing and the brain-damaged have in perceiving speech. It creates possibilities that would not otherwise have existed for research into aspects of the biology of speech and the presence (or absence) of the appropriate biological predispositions in infants and children. And, most generally, it enlightens us about the organic connection of speech to language.

Tactile Communication for the Blind

Although braille has been used since the 19th century as the primary reading method for the visually handicapped, it requires several years of careful schooling in its use and demands that a publishing organization exist to convert printed matter to braille type. Moreover, fewer than 20 percent of the sightless population actually acquire a truly literate skill in braille, i.e., read the equivalent of four or more books a year. The skilled blind person is therefore dependent on the publications efforts of an industry having limited resources for all reading matter.

Far more satisfactory would be an aid for the blind that permits the sightless person to "read" ordinary ink print by substituting the sense of touch for the missing visual sense. Were such a direct ink-print reading system available, the sightless person would be much more independent, perhaps able to read even the labels on prescription bottles or phonograph records. One such device is called the Optacon, a contraction of optical to tactile converter.

The story of the development of the Optacon illustrates the needed interplay of knowledge, intellectual curiosity, and inventiveness. Success with the use of the tactile sense as a substitute sensory channel for the handicapped had been only moderate by the middle 1950s. At about that time, Geldard and a number of his students had well under way a program of basic research in skin sensitivity that reached a watershed with the construction of a formal skin language called vibratese. Devised as a simple alphabetic code, people could process it at rates of up to 38 words per minute after a relatively brief training period. This research effort and its publication accelerated the efforts throughout the world to use the skin as a substitute for sight.

At the same time, a program was under way at the Massachusetts Institute of Technology for developing and improving sensory aids for the handicapped. One of its engineers, J. Bliss, moved to Stanford University and there engaged in collaborative research with J. Linvill, whose daughter was severely visually handicapped. The availability of a new technology for electronic circuit fabrication emerged in the early 1960s along with an extremely small and efficient device for transducing electrical signals to mechanical motion. With these and other technical advances, which permitted the design of a very compact processing device, along with the basic psychophysical information found in the current literature on cutaneous sensitivity, Bliss and his colleagues conducted a series of research studies with support from various government agencies. Based on the skin's ability to decode pattern information in real time, they arrived at specifications for a feasible processing device.

The program of research, which took about a decade for its fruition, culminated in the production of a prototype of the Optacon, which was tested by Linvill's daughter under various conditions and with a variety of print styles and sizes. She became one of the first blind persons in the world to read ordinary printed matter at speeds of up to 80 words per minute.

Hearing Aids

Despite the personal suffering and the loss of productivity arising from hearing impairments, and despite the great advances in technology that permit the realization of almost any signal processing scheme in a cosmetically acceptable hearing aid, present-day hearing aids are still quite inadequate for a large fraction of those with impairment. Roughly speaking, current aids provide nothing more than amplitude amplification that varies with frequency. This is adequate for the impairment, called conductive loss, that is characterized by poor conduction of the acoustic energy into the neural code. Such frequency-dependent amplification is not, however, at all adequate for impairments that involve malfunction of the sensorineural processes of the inner ear or auditory nervous system, the so-called sensorineural losses. These losses result in a degraded ability to resolve different acoustic stimuli, an increased susceptibility to background noise and reverberation, abnormal changes in the loudness of a sound with change of intensity or of duration, various types of sound distortions, and internally generated sounds. When these are the symptoms, no simple frequency-dependent amplification restores normal or even functional hearing. In case of severe loss, amplification may make it possible for the impaired person to know someone is speaking but still not be able to understand the speech.

Our existing knowledge about impaired auditory perception is not adequate to permit us to characterize the nature of the aids needed to correct for sensorineural impairments. Current work on this topic should in due course lead either to the development of improved aids or to a very clear understanding why such aids cannot be developed with existing technology.

Many new signal-processing schemes are now being explored as aids for the sensorineural impairments. Amplitude compression--the systematic change of naturally occuring amplitudes to a different range of amplitudes--appears to be promising for individuals with a reduced range. Frequency compression is also being explored as an aid to people who can hear only a limited range of frequencies. In both schemes the central idea is simply to recode the original information so that it is presented in those regions of frequency and amplitude in which residual hearing exists. Other schemes explore the use of multiple microphones to simulate our binaural hearing and thus reduce the background interference. Such aids may be particularly useful for individuals with a reduced capacity to comprehend signals in complex acoustic environments. Even more elaborate aids based on automated speech-recognition or speech synthesis systems are currently being investigated.

We know that the intelligibility of speech can be improved for many impaired listeners by transforming the speech signal in ways other than frequency-dependent amplification. For example, speakers can learn to talk in special ways so that a particular impaired person is able to understand more of what is said. They neither speak more loudly nor more slowly, yet they are better understood. Studies are presently under way to try to understand how these improvements are achieved.

In general and independent of the type of impaired hearing, signal processing must be devised that enables such people to make considerably increased use of their residual hearing. The minimum criterion of success is the ability to understand speech. To design appropriate processing, it is clearly necessary to characterize in considerable detail the nature of the several types of impairment that occur, and this may very well require a substantial amount of psychophysical and physiological research. How much research is difficult to say, since we only partially understand what it is we are searching for. Success will mean both the systematic development of improved auditory prosthetics and a contribution to improved audiological diagnosis and to increased understanding of normal auditory functioning. The latter plays an important role in the improved design of high-fidelity music equipment, more efficient telephones, and better design of acoustic systems that must be used in adverse conditions, for example, military applications.

Tactile Communication for the Deaf

We end this list of applications with a final example--the use of the skin as a substitute for hearing--where success has not yet been great and we are still awaiting a breakthrough. But hope is high, and a number of investigators are pursuing a variety of different paths. We outline some of their hopes here.

Some hearing loss is so great that there is no residual hearing to draw upon. In this case, the only possibilities for restoring the function of hearing involve sensory substitution: Sound must be transformed from acoustic vibrations to either visual or tactile patterns that are then perceived.

Although it is possible to learn to understand speech with reasonable accuracy by observing the lips of the speaker (lipreading) or by placing a hand on the speaker's face (the Tadoma method employed by some deaf-blind people), both of these methods have serious limitations: Lipreading requires adequate lighting, Tadoma requires direct physical contact, and both require proximity of speaker and "listener." For these reasons, attempts at more satisfactory technology have and are being made.

The first attempt to develop a tactile aid for speech reception, more than 50 years ago, used a single vibrator, acting as a loudspeaker. Hardly more than a placebo effect was noted, i.e., the deaf persons responded to the interest of the experimenter by working harder and giving improved performance through greater effort, but better information processing was not achieved. The basic bottleneck is that the frequency range that the ear appreciates is some 40 times as wide as the tactile range of 5 to 500 Hz. Obviously, some rearrangement or recoding of the signal must be provided to give the skin the perceptual span processed by the ear.

One class of devices arranged for a set of electrical filters that would "fan out" the speech frequencies over the skin so that a particular site, when stimulated, always represented a particular frequency band. Again, this met with limited success.

Another possibility was to follow the lead pioneer by Békésy. In the course of his work on the mechanics of the inner ear, he analyzed the mechanical and hydrodynamic properties of the cochlea. He calculated the values of the mechanical constants of the membranes that supported the sensory receptors of the ear. From these computations he built a dimensional model of the cochlea that, he claimed, vibrated in the presence of sound energy in the same manner as do the receptive tissues of the inner ear. He was able to demonstrate this correspondence by having observers place their arms on the vibrating "membrane" of the mechanical model and feel the change in frequency, loudness, and location of "sound" as he manipulated these variables by electrical or acoustic means.

A prominent German investigator, W. D. Keidel, decided to use the Békésy model, which had been designed purely for the purpose of advancing the understanding of the hearing process, as a practical, tactile speech-analyzing device. Because the skin accepts only low-frequency sounds, the speech frequencies must first be reduced to the proper range for the model to work. This was done by tape-recording the speech and replaying it for the model at a very low rate, which not only reduces the frequency to the proper values but also stretches time so that one

word may take several seconds to feel. Nevertheless, a few subjects did learn to understand some speech thus translated, and the investigator was encouraged to search for ways of getting speech sounds processed by computer and presented to the model so that "listeners" could feel speech patterns as they were being generated. A computer program was automated to scale the speech frequencies down to range suitable for the skin and at the same time clipped and joined the speech segments smoothly to produce a low-frequency version of the frequency-time relations of speech that lasted no longer than the original. Only limited tests have been made thus far, but this approach still seems to encourage further exploration.

Some positive results have been reported by other research groups studying tactile speech displays, which incorporate a variety of approximations to the spectral analysis performed by the ear.

We still need to determine the ultimate limits of tactile sense for speech communication, to clarify the principles governing the effectiveness of tactile displays, and to develop practical aids capable of functioning at a distance in real-world environments. In addition to contributing to the development of prosthetics for the deaf and deaf-blind, research in this area will provide increased understanding of the tactile sense, of speech perception, of design principles for sensory display, and of sensory substitution and human plasticity.

SUMMARY

This paper has reviewed what has been learned about the human senses, together with the relevant cognitive and motor components, by applying the scientific principles of the discipline called psychophysics. This information has been acquired, mainly in the past 130 years, by adapting and applying the methods of physics to problems of human beings interacting with various physical environments, ranging from very simple energy changes to complex dynamic systems, such as flying an aircraft.

As a result of these studies we now have a detailed and often quantitative understanding of many perceptual phenomena. The human senses and their associated cognitive and motor components are not passive systems reflexively reacting to the incoming stimuli, but are active transformers and processors of the applied stimulus. A great deal is known about how many of these transformations operate and their effects on perceptions.

We have also dealt with how this knowledge has been used in an attempt to solve, or at least understand, a number of practical problems. Our emphasis on sensory impairment arose not only because of the humanitarian component, but also because the best way to demonstrate a detailed understanding of some system is to be able to repair or otherwise ameliorate a defect in that system.

Success measured in these terms has been only partial, but in many areas great promise is evident. Our examples have exposed the false dichotomy between applied and basic research. Rather, as the section heading of these examples implies, there is an interplay between basic knowledge and the information and understanding gained by the attempt to apply that knowledge to concrete problems.