

Seeing as thinking: an active theory of perception

Richard Gregory

Theories of perception – of what happens to bridge the extraordinary gap between sensory stimulation and our experience of external objects – have a long history, of astonishing variety. Speculation goes back to the beginning of recorded philosophy – and scientific work on perception escapes the philosophical questions and dilemmas only when it narrows inquiry by over-blinkering specialization. How we see remains essentially mysterious after a century of intensive experiment, by such a variety of scientists that aims and communication can be lost between them. An adequate theory should include not only the favoured sense of sight but also: hearing, touch, hot and cold, taste, smell, balance and position of the limbs, the various kinds of pain; and tickle, from its irritation to sensuous pleasure and delirious laugh – making.

To the philosopher and the experimental scientist, it is how we see that offers the most exciting

questions, with hearing the runner-up, for sight dominates by its giving us immediate external reality. By simply looking we seem to understand what we see. This close association between seeing and knowing makes the sense of vision attractive not only to philosophers but also to experimental psychologists and physiologists who hope to discover in the brain mechanisms serving our experience and knowledge of the world. By coming to understand how we see might we not at one stroke also discover how we think, remember, formulate hypotheses, appreciate beauty and – most mysterious – accept pictures and words as symbols conveying not merely present reality but other realities distant in space and time? And if seeing involves all this, surely the net of understanding must be cast wide.

Perceptual theories form a spectrum -- from *passive* to *active* theories. Passive theories suppose that perception is essentially cameralike, conveying selected aspects of objects quite directly, as though the eyes and brain are undistorting windows. The baby, it is supposed, comes to see not by using cues and hints to infer the world of objects from sensory data but by selecting useful features of

objects available to it directly; without effort, information processing or inference. Active theories, taking a very different view, suppose that perceptions are constructed, by complex brain processes, from fleeting fragmentary scraps of data signalled by the senses and drawn from the brain's memory banks - themselves constructions from snippets from the past. On this view, normal everyday perceptions are not selections of reality but are rather imaginative constructions - fictions-based (as indeed is science fiction also) more on the stored past than on the present. On this view all perceptions are essentially fictions: fictions based on past experience selected by present sensory data. Here we should not equate "fiction" with "false". Even the most fanciful fiction as written is very largely true. or we would not understand it. Fictional characters in novels generally have the right number of heads, noses and even many of the opinions of people we know. Science fiction characters may have green hair and an exoskeleton-but is this novelty not a mere reshuffling of the pack of our experiences? It is doubtful if a new "card", suddenly introduced could be meaningfully described or seen.

The passive paradigm may, at least initially, seem more

acceptable as a scientific theory. It fits well with - and indeed essentially is -- the familiar "stimulus/response" notion in which behaviour is described as controlled directly by prevailing conditions. This is also familiar in engineering : in most devices input directly controls output ; and much emphasis is put on measuring input and output, and relating them by transfer functions or something equivalent, to describe the system. B. F. Skinner in his behaviourism claims to do much the same -- to give at least a statistical account of the relationship between stimulus (input) and behaviour (output) in animals and men. An engineer would go on to suggest "models", of what the internal mechanisms might be which transform inputs into the outputs. But, rather curiously. Skinner does not attempt to make this further step, and apparently distrusts it. He says remarkably little about brains, and at times denies memory and indeed all internal processes. His description is purely in terms of input output relations, with emphasis on how the probability of certain kinds of behaviour is changed by environmental-changes, especially "reinforcers".

Skinner himself has little interest specifically in perception, but passive theories of perception are in many ways similar. They

have the same initial scientific credibility, but are (I believe) essentially incorrect. They deny that perception is an active combining of features stored from the past, building and selecting hypotheses of what is indicated by sensory data. On the active account we regard perceptions as essentially fictional. Though generally predictive, and so essentially correct, cognitive fictions may be wrong to drive us into error. On this active view, both veridical (correct-predictive) and illusory (false-predictive) perceptions are equally fictions. To perceive is to read the present in terms of the past to predict and control the future. This account is very different from the passive story implied by Skinner's behaviourism, and most ably propounded by James J. Gibson and Eleanor Gibson (whose article is on page 711).

Why should one want to push all this stuff about "brain fictions" (as I do) when stimuli and responses are so easily observed, and so like the usual stuff of science? The essential reason is (I believe) very easily demonstrated, by common observation and by experiment. Current sensory data (or stimuli) are simply not adequate directly to control behaviour in familiar situations. Behaviour may continue through

quite long gaps in sensory data, and remain appropriate though there is no sensory input. But how can "output" be controlled by "input" when there is no input? The fact is that sensory inputs are not continuously required or available, and so we cannot be dealing with a pure input-output system. Further, when we consider any common action, such as placing a book on a table (a favourite example of philosophers) we cannot test from retinal images the table's solidity and general book-supporting capabilities. In engineering terminology, we cannot monitor directly the most important characteristics of objects which must be known for behaviour to be appropriate. This implies that these characteristics are inferred, from the past. The other highly suggestive--indeed dominating--fact is that perception is predictive. In skills, there may be zero delay between sensory input and behaviour. But how could there be zero delay, except by acting upon a predictive hypothesis? (Surely J. J. Gibson's description of perceptions as selections from the available "ambient array" will not do: it would have to be a selection from a *future* "ambient array" for the passive account to work: but this evokes a metaphysics we cannot welcome. The significance of prediction in perception has been

for too long almost totally ignored.)

It is the fact that behaviour does not need continuous, directly appropriate sensory data that forces upon us the notion of inference from available sensory and brain-stored data. This account is very much in the tradition of the polymath nineteenth-century physicist and physiologist, Hermann von Helmholtz, who described perceptions as "unconscious inferences". This notion was unpalatable to later generations of psychologists, who were over-influenced by philosophers in their role – sometimes useful, but in this case disastrous – of guardians of semantic inertia: objecting to inference without consciousness. But with further data on animal perception, and computers capable of inference, this essentially semantic inhibition has gone. Curiously, though, the kinds of inference required for perception are remarkably difficult to compute.

The recent engineering – science of Machine Intelligence is finding heavy weather designing computer programs to identify objects from television camera pictures. The reason seems to be (apart from the very large and fast computers required to perform the

operations serially) that the computer requires a vast amount of stored data of common object properties with ready and rapid access. It requires, in short, what we have called "fictions" to augment and make use of data monitored from the world by its camera eye, and – in machines dealing with real objects – its touch probes. In short: we may think of perception as an engineering problem, but it is a highly atypical problem even for advanced computer engineering, and it requires a special philosophy which is unfamiliar in science, because only brains and to a limited extent computers are cognitive.

The notion that interpreting objects from patterns is a "passive" business must strike the computer programmer engaged on this problem, in Machine Intelligence, as an extremely unfunny joke. His problem is to devise active programs adequate even for perceptual problems solved by simple creatures long before man came on the scene.

The notion of perceptions as predictive hypotheses going beyond available data is alien and suspect to many physiologists. Cognitive concepts appear unnecessary, even metaphysical – to be explained away by

physiological data. Certainly more physiological data are needed: but will they tell us by what mechanisms the brain's hypotheses are mediated, or will the "brain fiction" notion drop out as unnecessary? Prediction is dangerous, but there are surely strong reasons for believing cognitive concepts should be required for brain research, because the brain is unique, in nature, as an information handling system. (Or at least it is on an active theory of brain function.) With the development of computers, we now have other information handling systems to consider: it is interesting to note that to describe computers, "software" concepts are adopted, similar to cognitive concepts. More basically, what are essentially cognitive concepts are very familiar in all the sciences, but hidden under a different guise – the *method* of science.

Generalizations and hypotheses are vital to organized science, for the same reasons they are essential for brains handling data in terms of external objects. Science is itself not "passive" in our sense, but puts up hypotheses for testing, and acts on hypotheses rather than directly on available data. Scientific observations have little or no power without related generalizations and hypotheses.

Cognitive concepts are surely not alien to science, when seen as the brain's (relatively crude) strategies for discovering the world from limited data-which is very much the basic problem of all science. Scientific observations without hypotheses are surely as powerless as an eye without a brain's ability to relate data to possible realities-effectively blind.

The full power of human brain fiction is apparent when we consider how little current sensory information is needed, or is available, in typical situations. Here we do not need initially to consider particular experiments -- and indeed the intentional simplifications and restrictions of the laboratory environment can make the point less obvious - that behaviour is generally appropriate to features of the world which are not continually available to the senses. When you trust your weight to the floor, or your mouth to the spoonful of food, you have not monitored the ground's strength or the food's palatability : you have acted on trust, on the basis of the past. You have acted according to probabilities, based on generalizations from past events -and neither generalizations nor probabilities exist, except in your brain, for they are not properties of the world. Now suppose that you gave up acting on

informed guesses and demanded, continuously, direct selections of reality. How would you get on? Would you not avoid mistakes never fall through rotten floorboards, never be upset by bad food -never be misled by going beyond the evidence? Yes, indeed, if there were sufficient evidence available. But the fact is that there is frequently no possibility, or time, for testing floorboards or food. They must be taken on trust-trust based on the past as stored in the brain.

We have arrived at questions which may be answered by experiment. We can measure performance, in the partial or total absence of sensory data, and establish whether and how far perception and behaviour continue to remain appropriate. We find that we can continue to drive or walk, or perform laboratory eye-hand tracking experiments, through gaps in sensory data: and not merely inertially, for we can make decisions and change our actions appropriately during datagaps. We must then be relying on internal data. This requires an internal fiction of the world-which in unusual situations may be false. If the situation is unfamiliar, or changes in unpredictable ways, then we should expect systematic errors generated by false predictions. Errors and illusions

thus have great importance for active theorists: they become obsessively used tools for discovering the underlying assumptions and strategies of the perceptual "computer" by which we infer - not always correctly - external objects from sensory data.

Looking at books written by passive and active theorists, we find an amusing difference between their indexes. Passive books devote much space to stimulus patterns, but very little to the phenomena of perception: spontaneous reversals in depth, changes into other objects, distortions, perceptual paradoxes in which the mind reels by being apparently confronted by logically impossible objects. Active theorists fill their books with examples of such phenomena, interpreting them in various ways, while the passive theorist ignores them, or writes them off as too trivial to concern him. But neither uncertainty nor ambiguity, neither distortion nor paradox, can be properties of objects: so how can we *perceive* uncertainties, ambiguities distortions or paradoxes if perception is but a passive acceptance of reality? This simple though surely powerful argument is not raised or answered by passive theorists. By playing down the obvious phenomena of perception (such as

illusions, found as children's puzzles) passive books may look academically safe – but at the cost of leaving out what is most interesting.

We may now return to the point that, although we regard brain function as physical, physical and engineering concepts are not adequate for describing some aspects – especially perception of objects. This only appears to be a metaphysical statement if an extreme reductionist view of science is adopted. This matter is controversial: there are eminent scientists who hold that knowledge of a hydrogen atom and the laws of quantum mechanics are sufficient to describe, in principle, any physical situation. Others hold that even common effects such as friction, heat, inertia or gravity (let alone brain function) could not in principle be described in these elementary terms. They hold that with increased complexity, and organization new properties arise requiring new concepts to describe them. It would certainly be difficult to ascribe the notion of “cognitive fictions” to a hydrogen atom! (But it would be equally difficult to ascribe such concepts as servo-control, or even image – forming – so this is not a special objection to the “cognitive fiction” notion.)

There is a strong reason (apart from consciousness) why we wish to separate descriptions of aspects of brain function from physics. This is however a very tricky problem, easy to over – state and to misunderstand. Granted that brain activity is physical, we wish to hold that brain states representing information and problem – solving are not usefully described in terms of physical restraints. Consider the black marks (letters) on this page. They are physical (ink absorbed by paper), but their arrangement, surely, is not to be understood by the principles of physics. For this we must call upon English spelling and grammar and upon the structure of what I am trying to say. In the vital respect of their order, they are free of the ink and paper of which they are made. If their order were determined directly by their material and its physical properties (as in crystal structure) then they could not serve as symbols. Being in this sense free of physical restraints, and given receptive brains (or computers) then they can serve as symbols: to represent objects in other time and space; or abstractions which do not exist, in the sense that objects exist. This is true for all symbols: pictures, words, mathematical and musical notations, video and audio tapes, computer tapes. But symbols are

powerless (or are just like any other objects) in the absence of brains or other information – handling systems. Evidently symbols must affect brains in some more or less lawful manner: but for this to be possible the relevant brain states must – like the typist's or compositor's characters – be free to adopt information storing and representing orders. So *they* must in this rather limited sense be free of physical restraint, though not quite isolated from the rest of the physical world for learning and perceiving to be possible.

The celebrated (and I believe essentially misleading) Gestalt theory of perception postulated physiological restraints to explain many visual phenomena, such as preference for, and distortion towards, figures of “simple” and “closed” form. Visual forms were supposed to be represented in the brain by similarly shaped electrical brain fields – circles by circular brain traces, presumably houses by house – shaped brain traces. These brain traces were supposed to tend to form simple and closed shapes, because of their physical properties; much as bubbles tend to become spheres, as this form has minimum potential energy. Now this implies that visual “organizations” and distortions are due to physical restraints and forces which will not in general be

relevant to the logical problems the brain must solve to infer objects from sensory patterns and stored data. This is quite different from a cognitive account of perceptual distortions, and other phenomena which may be supposed to arise from misapplication of strategies quite apart from the physiology involved. Using a slide-rule, an error may be due to physical errors in the rule itself, or to misapplication of the rule for the problem in hand. This is exactly the distinction involved here, between physiological and cognitive errors.

We should expect physiological restraints to produce the same effects for any object situation (for example after – images, due to retinal fatigue, to any bright light). Misplaced strategy errors should, on the other hand, be related to the kind of perceptual inference, from sensory pattern to object, being carried out. So the point is that the physiology should only produce errors when it is exerting *general* restraints. We should not expect this except in abnormal situations, such as when the physiological “components” are driven beyond their dynamic range. Considering phenomena of perception, such as ambiguous, distorting or paradoxical figures: do these figures upset the physiology, or select inappropriate strategies, to

generate errors? In these cases, it seems to be the object significance of the figures which is relevant. So these phenomena seem quite unlike after-images – here it is not so much the physiology as the cognitive strategies which we need to discover. This needs a different (but still a “scientific”) way of thinking, and powerful experimental techniques, to discover cognitive strategies and how they can mislead.

To separate errors due to physiological restraints from errors due to misplaced strategies surely has importance beyond understanding perceptual errors. The same distinction (between physiological and cognitive processes, and how either can go wrong) might be important for understanding mental illness. If schizophrenia is errors in the brain’s strategies for developing hypotheses of external states of affairs, this should be understood not only in terms of biochemistry and physiology but also in terms of the strategies by which we normally cope with things. Perhaps this matter of strategies is hidden by the apparent ease with which we continually solve problems of the utmost difficulty to computer programmers: and which receive false answers when their programs are inappropriate. Seeing a table as something to

support a book upon is to solve a problem so difficult it challenges the most advanced computer technology, and yet to us it is so simple that a passive theory of perception may seem plausible. This shows that passive theories may be so misleading as to hide aspects of brain function we must see clearly to understand not only perception but all mental processes and how they can go wrong.

Recent discoveries by physiologists, especially by electrical recording from single brain cells during controlled stimuli to the eyes, are so clearly important that they tend to dominate much current thinking about perception. The problem of how sensory patterns are interpreted in terms of objects tends to be ignored. The important physiological discovery is that certain stimulus patterns (lines of certain orientation, or movement, etc) produce repeatable activity in specific brain cells. This discovery came as an unpalatable shock to passive theorists who tend to ignore brain function. To active theorists, it gives a clue to the kinds of data accepted for building object-hypotheses. One might think from this that passive theories would drop out, leaving the field of physiologists and active cognitive psychologists to work together in blissful harmony.

Actually things are not quite like this : the physiological advance is so concrete, and clearly important, that many physiologists and cognitive psychologists feel that finding more feature analysers, and more abstract analysers, is the sole path we need to follow to understand vision. But is it? The physiological mechanisms being discovered relate to stimulus patterns only, and not recognition of objects its hypotheses. The physiological account thus remains passive, and so essentially inadequate, for the same reasons that cognitive passive accounts are inadequate.

The task ahead is to relate physiological processes not only to direct input-output links, as in reflexes, but also to the brain's logical and correlating activity endowing it with the power to predict. This will require further physiological data, and current techniques are providing extremely important new information so this will surely be available. Experiments on the phenomena of perception itself, in animals and in men – essentially on how patterns are interpreted as objects -- has confusions (or at least impeding disagreements) in its philosophy, and a lack of powerful research techniques. Some of the most interesting clues are at present coming from studies of

development of perception in babies. Early changes of the nervous system as a result of experience are now being discovered, which will perhaps help to tie up, or relate, physiology and cognition. Possibly the most fundamental and rigorous ideas are coming not from biology but from attempts to program computers to see and handle object-relations. It proves necessary to make the computers develop hypotheses and select the most likely, given the data from its glass eye.

There is more to this, for some computer programs designed to give "scene analysis" (recognizing objects from pictures by computer) assign alternative object probabilities to selected features in the picture: and then change these probabilities, according to probabilities assigned to other features of the scene. For example, a given shape may be a box or a building. If what is taken to be a hand is above it, then the probability of the box hypothesis will be increased and the building hypothesis decreased-for hands are generally too small and too low to be above buildings, but not above boxes. Now this gives interactions, due to conditional probabilities, which may generate visual effects in computers or brains quite like the old Gestalt phenomena, but for an entirely different reason. The

reason is to be understood in terms of cognitive strategies or procedures for making effective use of data for deciding what objects are present in the scene.

In Machine Intelligence only precisely formulated theories are adequate: any gaps or errors in the theory show up as errors in the machine. At present machines perform only the simplest tasks, and are easily confused by shadows or small changes we scarcely notice.

Although the difficulties in Machine Intelligence demonstrate all too well how little we know, it now seems that we are beginning to understand ourselves – the inference – mechanisms of our humanity – by inventing adequate concepts for machines to infer objects from data: to perceive our world with their metal brains and human-devised programs. Is this science fiction? Yes-but like all fiction it may be largely true.

Philosophically, this is not the end of the matter. Behaviourism, with its related passive theories of perception, is unconcerned with what goes on between the senses and behaviour: indeed denies that anything goes on. This may be a legitimate expedient for focusing attention upon certain questions in behavioural research; but as a philosophy it is a kind of nihilism with a built-in contradiction. We

are supposed to accept the behaviourist's writings as expressing his observations, thoughts and judgments: which in these same writings he denies having. We are reminded of the poignant postcard received by Bertrand Russell saying: "I am a solipsist-why are there no other philosophers like me?"

*Professor Gregory is head of the
Brain Perception Laboratory,
University of Bristol.*