

The Cognitive Architecture of Perception

Juan Vázquez



The Cognitive Architecture of Perception



Juan Vázquez

Edição



Apoio



This is a digital, open-access version of the print-format edition published originally by the Universidade de Porto in 2014 with ISBN 978-989-8648-22-8

FICHA TÉCNICA

Título:

The Cognitive Architecture of Perception

Coleção:

MLAG Discussion Papers, vol. 7

Direção:

Sofia Miguens e João Alberto Pinto

Autor:

Juan Vázquez, Universidade de Santiago de Compostela

DOI: <http://dx.doi.org/10.15304/op.2016.956>



This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>

Foreword	4
Introduction	7
I. Stimuli and stimulation	11
II. The architecture of the perceptual system	17
III. The categorial framework	27
IV. Subcategorial and categorial perceptual information	49
V. Categorial perception	63
VI. The presentational nature of perceptual content	73
VII. Categorial perception and the empirical justification of beliefs	83
VIII. Categorial perception and scepticism	91
Bibliography	101

Foreword

According to Hayek, to perceive is to classify the world in accordance with the properties that the objects within it share with others we have experienced in the past. The author of this book refers to this process as “categorization”, thereby emphasizing the hierarchical organization of the perceptual content of memory. Perception is indeed an interpretation of the world in which we live, in accord with our memory of it. As Helmholtz in effect pointed out, we do not just remember what we perceive; we also perceive what we remember. That said, it is important to add that the properties shared by perceived and remembered objects are not simple physical attributes such as colour, shape or size, but also the relationships among them. As in *Gestalt* psychology, what defines an object in perception is the relationships among its elements. A rectangle is defined by the relationships among four angles; a rose by the relationships among the shape of its petals, its aroma, and its colour. The perceptual code is a relational code, and relationships are the essential attributes of any known, recognized or perceived object.

It is with these premises, as they relate to their cerebral substrate, that Dr. Vázquez approaches perception. In the light of modern cognitive neuroscience, the system that classifies - “categorizes” - perceptions consists of an assembly of distributed, overlapping and interconnected cortical neuronal networks. These constitute the nervous substrate not only of perception but also of memory and the other cognitive functions. These networks, which I call *cognits*, are the irreducible units of memory and knowledge, formed by synaptic reinforcement among the neural assemblies that represent the elements of sensation that are simultaneously received from an object or situation. The cognit, like the network that forms it, is thus defined by relationships (connections). It is worth pointing out here that although perception has mainly been studied in the visual world, with regard to its retinotopy, cognits map everywhere topologically, not topographically. That is, they represent relationships – and not necessarily spatial relationships – among their elements. They are scale-free, like matrices of rubber bands linked together by multiple nodes; you can stretch, twist, squeeze and bend such a matrix without altering the relationships among its nodes (the associated groups or assemblies of neurons). This is how cognits are structured in the cortex. For this reason, in psychophysics, perception is a relational phenomenon: it is relational by virtue of the relationships that define

the cognit and its context. It is these relationships that explain the visual illusions mentioned in this book.

The reader may think that with these remarks on Dr. Vázquez's model I am just touting my own ideas. Not so. My aim is to highlight the conceptual significance of his essay so as to make it applicable to a wider field of cognition than that to which he judiciously limits it. Here I take the liberty of broadening its scope, without modifying the essence of its arguments, to argue that sensation, categorial perception and conceptual semantics all belong to one and the same continuum. And it is here, in its hierarchical aspect, that the author's concept of perceptual category is especially apt, since categories are hierarchically organized, in the brain and in the mind. It is impossible to appreciate Dr. Vázquez's model in all its depth without considering it in relation to the hierarchical structure of the cerebral cortex. This structure, shaped by the experiences of the organism in the course of phylogenetic and ontogenetic development, comprises sensory structures in primary sensory cortex, at the lowest level of the cognitive hierarchy, and associative areas in parietal, temporal and prefrontal regions at the higher levels.

In principle, according to our current understanding, perceptual categorial networks are formed through the synaptic mechanisms of association that act at all levels of the cortical cognitive hierarchy. Moreover, since at any level a cognit is formed from lower, more specific cognits (the category "table" emerges from the large variety of objects to which the word is applied), lower categories are nested within higher categories that are not only more abstract but also more complex, in that, while conserving the same basic relational structure that makes them categories, they embrace multiples versions of the same object.

Due to emotional content, learning, disuse or aging, the cognitive networks of the cortex are permanently in conditions of unstable equilibrium. Synapses are lost or reinforced in time, in ways that are imponderable by current methods; this is one more reason to speak, as the author does, of perceptual presentation rather than perceptual representation. Nothing is perceived or remembered exactly like the first time. Behind the subreptitious modifications of cognitive neural networks lies one of the causes of the false memories with which judges and juries must reckon.

In the course of behaviour, speaking, or reasoning, as one of these activities is pursued with intent and with attention to its aims, perceptual cognits are brought into play and take their place in the operations of the perception-action cycle. This cycle is just the formalization in the cerebral cortex of the cybernetic loop that governs the relationship of the organism with its environment at all levels of the central nervous system. It is in this context, as the author notes, that categorial perception acquires dynamical significance, that is to say, when it enters the realm of consciousness, attention and working memory. And according

to recent work in cognitive science it is in this “presentational” framework of perception that the blending of semantic and categorial contents of perception is consummated in the highest levels of the cortical hierarchy by which it is supported.

Joaquín M. Fuster *
Los Angeles, California



* Professor of Psychiatry and Biobehavioral Sciences, UCLA - Semel Institute for neurosciences and Human Behavior - School of Medicine

Introduction

My aim in this study or essay has been to provide an explanation of the cognitive processes involved in perception, making use of the information that is currently available thanks to neurophysiology, cognitive psychology and anthropology. My main concern has therefore been the analysis of those types of perception that involve knowledge - in other words, those that can be employed in the empirical evaluation of knowledge.

Perception is a source of information, an informational mechanism serving a living organism. The information acquired by the organism from its environment may accordingly depend to a large extent on its circumstances, purposes, modes of action, and practical needs in general; and the kinds of perception through which the organism obtains this information will be similarly diverse. If the activity pursued by the organism is motion through its environment, or interaction with it, or the paying of attention to one of its multiple aspects, then it is pointless for it to attain explicit awareness of each and every one of the stimuli it receives, regardless of their relevance to that particular activity; to do so would involve useless effort and distract from the activity pursued. Imagine, for example, a pedestrian who wants to cross a street using a zebra crossing: for this purpose it will suffice to detect the zebra crossing and, upon arrival thereat, ascertain whether there are any vehicles approaching from either direction. All other information about the zebra crossing (the width of its stripes, their colour, the state of the paint, etc.), or about a multitude of other features of the pedestrian's visual field, is irrelevant to the proposed action, and it would therefore be pointless to pay these features any attention. And if they are not paid any attention, they will not be perceived - not, at least, with the clarity required for their identification. Indeed, for an activity such as crossing a street, it suffices for perception to be of an almost mechanical nature, with no greater precision than is appropriate to the activity in question. There is no need here for the perceptual experience to involve the subject's conscious awareness that he or she is paying attention to anything in particular; the information acquired is for practical use, an ancillary to action.

In cases such as the situation described above, perception guides our actions; it is a rich source of information, but the information needed for effective performance of the proposed actions does not need to include details of the environment that are irrelevant for their guidance. Only occasionally do we

pay attention to some object, property or event in our environment with the purpose of achieving its conscious perceptual identification. This occurs, for instance, when in perception we seek justification for some belief; when, for example, I look out of the window to check that the plum tree in my garden is in flower. But even in such cases as these, attention is selective, and among all the stimuli provided by the environment it will accept only those that are relevant to identifying the state of the plum tree as “in flower”. An enormous quantity of additional information that is available in the stimuli provided by the plum tree (the exact arrangement of its branches, the number of flowers that are open, their exact colour, etc.) is irrelevant to this particular identification task and will accordingly not be taken into account by selective attention, with the result that these traits will not be perceived - not, at least, with the clarity and sharpness proper to conscious perceptual identification. It is conscious perceptual identification that I will be concerned with here, because it is this kind of perception that is of special cognitive interest: only in regard to this kind of perception does it make sense to ask whether perception can or cannot justify beliefs.


There are many other kinds of perception encompassing perhaps the immense majority of our ordinary perceptions for which inquiry into their epistemic value is meaningless, and is something that we never in fact embark upon. The immense majority of our ordinary perceptual experience is used to guide our interaction with our environment, and accordingly involves the processing only of such aspects of external or internal stimuli as are relevant to the activity being pursued. Moreover, the perceptual reception of such stimuli need not be accompanied by explicit awareness of the information that is being used in the performance of this activity; indeed, explicit awareness, besides wasting energy, might even interfere with performance.

If for most of our ordinary perceptual experiences it makes little sense to ask about an associated epistemic value, such inquiries make even less sense in regard to perceptual experiences in which the subject not only has no explicit awareness of paying attention to anything in particular, but would, if questioned, profess awareness of not perceiving anything at all, as in cases of blindsight. And much the same may be said in regard to the various kinds of visual agnosia, in which the affected suffer no loss of sight, and can respond to what they see in certain practical ways, but without this perception of the objects seen being accompanied by their recognition, so that the perceptual information received must be regarded as lacking any epistemic or cognitive value.

In short, there are many different kinds of perceptual experience, and it is therefore often improper and inappropriate to speak of perception in general, as if there were only a single kind. To avoid multiple misunderstandings, I wish to make it clear that I am concerned here basically with the kind of perceptual

experience in which the perceiving subject consciously identifies an item in his or her visual field, since it is in relation to this kind of perceptual experience that it makes sense to ask whether perception can or cannot justify beliefs.

With the object of my analysis thus specified, I now point out that its study requires consideration of at least three types of elements or constituents: the *stimuli* provided by the environment or by the subject's own organism; the *architecture of the perceptual system*; and the prior *knowledge*, or *categorical system*, that the subject has acquired in the course of his or her development. Schematically:




Lest I be misunderstood, I stress that, as will emerge in the corresponding chapters below, phenomenal content (sensation) and conscious categorical content (categorical perception) arise simultaneously and are mutually interdependent - which is not to say that they cannot be distinguished experimentally or that, prior to categorical identification, there is no sensory basis that both determines that identification and is categorially determined by processes of selective attention. What I pursue in this study is the analysis of each of the above three components and the determination of their separate and joint roles in the processes of perception, together with the consideration of a few applications.

As a rough guide to the reader, the contents of the chapters of this book may briefly be described as follows. Chapter I is devoted to a discussion of stimuli and stimulation, since the use to be made of these concepts is somewhat different from their customary usage in ordinary language or research. Chapter II is a brief sketch of the architecture of the human visual system, with emphasis on factors with decisive influence on the processing of stimuli. The reason for concentrating on the visual system is that this is the sense system that has been most intensively studied and is best understood. Nevertheless, many of the conclusions reached concerning sight are applicable to the other senses. Chapter III presents the genesis, constitution and development of category frameworks, and their important role in cognitive development, in the light of experimental findings from cognitive psychology and anthropology. Although the terms "category" and "concept" are often treated as synonyms in the cognitive science field, in this chapter I follow Gregory Murphy (2004) in distinguishing between categories and concepts, and this distinction will play an important role in the analysis of perceptual content. In Chapter IV, following Thomas

Metzinger (2003), I suggest that it is necessary to distinguish between categorial and subcategorial perceptual information or content; this distinction allows one to deal with many of the difficulties faced by both perceptual conceptualists and perceptual non-conceptualists. Once the distinction between categorial and subcategorial information has been made, and the presence of both these types of content in perceptual processes has been established, Chapter V discusses the configuration of categorial contents, and the important cognitive role they play in perception. The discussion of perception is completed in Chapter VI, which argues that perceptual content, whether categorial or subcategorial, is presentational rather than representational. This position, though apparently close to direct realism and disjunctivism, differs in its epistemological setting. It is this different setting that allows me to argue that perceptual content is presentational, and therefore that perception affords direct access to the world, while avoiding the pitfalls that beset both direct realism and disjunctivism.

The next two chapters show how the theory of perception developed in Chapters I-VI provides a criterion of truth for perceptual statements and a refutation of scepticism. The criterion of truth presented in Chapter VII is close to the correspondence criterion, but avoids the difficulties of the classical view of truth as correspondence. In relation to the proposed truth criterion it is shown that the beliefs expressed by perceptual statements receive empirical validation in the categorial contents of perception. Finally, in Chapter VIII, I refute epistemological scepticism, the thesis that through perception we can know nothing of the world around us. The refutation proceeds by exposing the two highly problematical premises on which the sceptical thesis is founded: the first (possibly adduced by the sceptic only as a working hypothesis), that there exists a world “in itself” that is independent and distinct from the world that is experienced; the second, a logical consequence of the first, that the content of perception is representational. If the account of perception developed during the first six chapters of this book is correct, these two premises on which the sceptic relies are groundless, and scepticism is accordingly unjustified.



I. Stimuli and stimulation



In the psychology of perception it is customary to distinguish between *distal* and *proximal* stimuli. A distal stimulus is any entity or property of the world that is or may be an object of perception, such as an apple, a chair, a colour, or a shape. The epithet “distal” refers to the stimulus generally being located at a certain distance from the sense receptors, as in vision or audition. In contrast, it is called proximal stimuli to the different types of energy (photons, sound waves, molecules, forces, etc.) that, once emitted or originated by the distal stimuli, directly activate the specialized receptor cells which in the course of evolution have developed to support the senses of sight, hearing, etc., and which respond to these agents by sending electrical signals to the brain. From this point of view, it appears logical to consider stimuli as the causes that trigger our perceptual experiences, distal stimuli being remote or indirect causes and proximal stimuli direct or immediate causes.

There is nothing objectionable about the above uses of the term “stimulus” if we are interested in studying or explaining aspects of perceptual capacity such as visual acuity. However, if our objective is to explain perception itself, this view of stimuli is less than satisfactory. Let us first consider distal stimuli. The objects and properties of the world, to which this term is applied, are in fact not so much the cause of perception as its ultimate effect, the result of the processing to which our brain subjects the corresponding proximal stimuli. When we see an object in our environment, the object perceptually experienced as external, as “out there”, is the object as it is identified by a human cognitive being in the process of perception. But this identification is the result of the processing to which our brain subjects the proximal stimuli that directly affect the sensory receptors and that these receptors transform, through the phenomenon of *transduction*, in electric current or action potentials. It is in this form, that the information emanating from the objects of the world arrives at the brain and that the brain interprets in terms of objects of the world that are independent of the perceiving subject; but we are actually dealing with a world of objects and properties generated through the corresponding processes of perception. In short, what experimentalists consider as distal stimuli are not in fact stimuli at all, since they are not really the cause but the ultimate effect of a process: the perceptually identified entities and properties generated by the perceptual processes set afoot in a normal individual when his or her sensory receptors are activated by the stimuli emanating from those putative entities.

To say, for example, that the PC on which I am writing is the cause of my perception of it can only make sense if in the process of perception I could

become, as Putnam would say, the Divine Eye, with access both to the PC as it supposedly is “in itself” and to the PC as it is experienced in the process of perception. If that were the case, then I could indeed say that the PC “in itself” is the cause of my perceptual experience of it; but in practice we have perceptual access only to the PC as it is experienced in the act of perception. And, as I argued above, the perceptually identified PC is not the cause of my perception of it but an ultimate effect, the perceptual identification resulting from the processing that my brain has performed on the stimuli emanating from that which I identify perceptually as my PC.

If it were a bat, for example, that perceptually identified the PC, we could be absolutely sure that the content of that perceptual identification would be significantly different from the content of the perceptual identification we human beings perform in regard to the “same” object. And this would not justify our saying or thinking that the perceptual identification of the PC performed by the bat is any less real than ours; more than one human being has been so unfortunate as to collide with a glass door, something a healthy bat would never do.

Nor can proximal stimuli be identified as the real sufficient cause of our perceptual experiences. In the first place, proximal stimuli are themselves very different from what human cognitive beings become conscious of in the process of perception. Whatever their wavelength, the bearers of light signals, photons or waves, are not themselves coloured; molecules, the bearers of olfactory signals, do not smell; and so on. Colours, shapes, movements, smells, sounds, tastes, etc., are only apprehended as the result of the perceptual processing to which our cognitive apparatus subjects those proximal stimuli. Again, the difference between one colour and another – a qualitative difference – is not the same thing as the purely quantitative difference between the wavelengths of the light waves that constitute the corresponding proximal stimuli. In fact, just like distal stimuli, proximal stimuli, to which the relevant scientific disciplines attribute objective independent reality, are the result of cognitive processing, in this case their processing, with the help of instruments of observation, by these scientific disciplines. Thus in the case of proximal stimuli, too, we are dealing with a pre-codified and conceptualized reality, the construction of which has already involved, directly or indirectly, acts of perception.

If what we want to do is to explain the content of our perceptual experiences in terms of a reality that can be identified without involving perception, then the only stimuli it makes sense to admit are non-conceptualized constituents of the world that in some way involve no prior perceptions. There is evidently nothing we can say about such constituents, and the only reason for postulating them is that normal perception involves the activation of our sensory receptors by a cause located in the external world or in our own organism, and that the

perceptual experience ceases with the cessation of the stimulation. If I am seeing some object and I close my eyes, I no longer see it; and if I am touching it with my hand, the tactile sensation also ceases as soon as I no longer touch it. In short, in an account of the process or processes of perception we must postulate that these processes have an external cause and that, strictly, it is to this external cause that we should attribute the function of stimulation, not to a pre-codified and conceptualized reality. As has already been explained, this does not mean that empirical studies of perceptual acuities and deficits, or of cerebral function, cannot legitimately continue to use the term “stimulus” in the two habitual senses described above; but to explain the phenomenon of perception itself we cannot accept as its external cause a pre-codified and conceptualized reality involving prior use of the processes of perception.

I repeat, what in the usual acceptation of the term is considered to be a distal or proximal stimulus cannot strictly be the cause originating perceptual experience, because both in the case of experimental scientific observation and in the case of perception, what is observed or perceived is a pre-constituted entity that belongs to the end of the observational or perceptual process, and that is therefore not its cause but its final effect. The true causes, the stimuli that originate our perceptual experiences, are to be found at the start of the process; and since all that we are conscious of in the processes of experimental observation and perception is the result of their processing, we know nothing of these stimuli *sensu stricto*, and can say very little about them other than to attribute to them the causal efficacy proper to any stimulus.

Accordingly, in the whole of this essay on perception the term “stimulus” will refer to *whatever it is that activates our sensory receptors*, and the term “stimulation” to *the fact that these sensory receptors are activated*. If asked about their nature, from a perceptual point of view I must reply that I know nothing about them, that they form no part of the content of our perceptual experience. And if anyone wishes to consider objects and other properties of the world as the remote cause of our perceptual experience, and to consider the various kinds of proximal stimulus postulated by diverse scientific disciplines as its proximal or immediate cause, I have no objection; for many research purposes it will be a satisfactory point of view since, however hard we try, we will never be able to escape from the codified world of our experience, regardless of whether its codification be performed by science or in the framework of ordinary knowledge. To use Donald Davidson’s expressive words, “we can’t outside our skins to find out what is causing the internal happenings of which we are aware” (Davidson, 1989, p. 312).

But if, as I maintain, we can know nothing of the putative uncoded stimuli that impinge on our sensory receptors, then why postulate their existence? Why is it necessary to postulate them? The answer is obvious. If we close our eyes,

the retinal receptors are no longer activated and, as a result, we cease to see what we were seeing before we closed them. Similarly, if we take our hand off the object we were touching, our tactile sensations cease; and the same goes for the other senses, *mutatis mutandis*. We must, therefore, postulate the existence of external agents that activate our sensory receptors and originate our perceptual experience. Which does not mean that in exceptional circumstances of perceptual abnormality there cannot occur perceptual experiences that have little or nothing to do with sensory stimulation, such as hallucinations.

Something else we can say about stimuli and stimulation is that when two or more subjects access the world through their various senses, the stimulation they receive must be, if not “the same”, then at least very similar, since in normal conditions the stimulation we receive leads us all to perform the same sensory discriminations in regard to each of the senses. If this were not so, our coinciding in our perceptual identifications of the same items would be inexplicable, as would the process of learning, that capacity of all normal subjects to learn to group the sensory information derived from the processing of stimulation in classes or types of things, events, properties, and other kinds of phenomena shared by the members of a community.





II. The architecture of the perceptual system



For convenience, and also because sight is the sense that has been most extensively studied and is best understood, I will concentrate here on the architecture of the visual system, though the results of my analysis will be broadly applicable to the other senses.

As we saw in Chapter I, the stimuli that activate the sensory receptors of the various senses are one thing, and the results of their processing, the “things” we become conscious of in acts of perception, are another. Between the reception of stimuli by sensory receptors and the final result there is a long process in which, in so far as visual perception is concerned, we may distinguish the roles of three components: *a*) the photoreceptor cells of the retina (the cones and rods); *b*) ganglion cells, the axons of which project to the brain via the optic nerve; and *c*) the primary and associative areas of the visual cortex, where the processing of visual stimulation occurs.

a) Retinal photoreceptors

The retina is located at the surface of the back of the eye, and basically comprises three layers of cells: a layer of photoreceptors (cones and rods), a layer of bipolar cells, and a layer of ganglion cells (Figure 1).

In the processes of visual perception the stimuli received from the outside world impinge directly on the photoreceptor layer. The photoreceptors, the only cells that are sensitive to light, communicate via synapses with the bipolar cells, and these latter synapse in turn with the ganglion cells. According to neuro - physiological estimates, each human retina contains about

126.5 million photoreceptor cells, of which 120 million are rods and 6.5 million cones. The cones allow us to see colours and distinguish shapes, and as we shall see below, it is basically the stimulation of cones that allows us to identify objects, events, properties, and so on. Rods, which are much more sensitive to light than cones, are responsible for vision at night and in dimly illuminated environments in general. Unlike cones, rods do not support colour vision or




Figure 1 Retinal Circuit

discriminate shapes, and provide relatively poor resolution; practically all they detect sharply is movement. The fovea, the central region of the retina, only contains cones, and the ratio of cones to rods decreases with distance from the fovea. When we wish to see anything sharply, we must fix our gaze on it so that the stimulation emanating therefrom will be received by the cones of the fovea (Figure 2).

These features of the architecture of the human visual system already have some important implications. One is that in the process of perception the retina does not act like a mirror. A mirror reflects alike everything placed before it; but the information passed on by the retina depends on which of its receptors are stimulated. If it is the foveal cones, the processing of these stimuli can produce a clear vision of whatever is in that part of the visual field. If, on the other hand, those “same” stimuli are received by the rods of the peripheral retina, their processing can only produce imprecise information regarding that part of the visual field. Thus the “same” stimuli can afford very different information depending on whether they are received by the cones of the fovea or the rods of the peripheral retina.

Consider an individual, S, working in his study; he may be sitting before his computer, or reading a book. The door of the study is located to his left or right, in his peripheral visual field. If the door is open and someone comes in, S will perceive a kind of shady presence moving in the doorway, because the rods of his peripheral retina will detect his visitor’s movement; but if S does not turn his head, and so prevents those stimuli from being received by his foveal cones, he cannot see whether his visitor is male or female, tall or short, or fat or thin. For him to obtain that kind of information, it is the foveal cones of his retina that must be stimulated.

Let us examine a second example. Suppose someone is driving along a road, on which she is accordingly concentrating her gaze, and that there is a dog lying by the roadside. Her passenger asks “Did you see that dog?”, and she replies that she did not. But her passenger, a philosopher, remarks “Yes you did, it’s just that you’re not conscious of having done so”. From a neurological point of view we would have to point out that the driver may very well not have seen the dog, even though it was in her visual field; the dog having remained immobile, and the stimuli emanating from it having been received by the rods of the driver’s peripheral retina, the processing of those stimuli will not have




Figure 2 Retina

allowed her to distinguish the colour and shape of the dog from its background, and she will thus have lacked two pieces of information that are fundamental for visual identification of something as a dog.

Thus the first conclusion I wish to draw from these neurophysiological facts is that, exposed to the same stimuli, different subjects, or the same subject at different times, can make different perceptual identifications depending on the type of receptors that are stimulated. If at one time someone fixes his gaze on the central area of Rubin's vase (Figure 3), what he sees clearly will be the colour and shape of this part of the figure; and if he has acquired the concept of a vase, he may see a vase of that colour and shape.



Figure 3 Rubin's Vase

But if he fixes his gaze on the periphery of the figure, what he sees clearly will be the colour and shape of that peripheral area, and he may also see - configured in that colour and shape - two faces looking at each other.

Two individuals can be looking at Rubin's vase from such closely adjacent positions that they receive practically the "same" stimulation from it, and yet, if one has fixed his gaze on the center of the figure and the other on the periphery, so that as a result of selective attention the stimuli received by the foveal cones of one are received by the rods of the peripheral retina of the other, and *vice versa*, then these two individuals will have different sensory and perceptual experiences, despite receiving practically the same stimulation. Which leads us to the general conclusion that, given the "same" or very similar stimuli, different subjects may have different sensory and perceptual experiences.

As was pointed out in Chapter I, sensory receptors – in this case the sensory receptors of the retina (cones and rods) – do not just receive stimulation from the outside world (or from the individual's own organism), but also transform it into electrical signals – so called action potentials - that are sent to the brain for processing. Let us examine this process.

b) Ganglion cells

The photoreceptors of the retina synapse with bipolar cells, which synapse in turn with ganglion cells. Retinal signals travel through the

optic nerve along the axons of the ganglion cells, and in the lateral geniculate nucleus of the thalamus are relayed to the primary visual cortex (Figure 4). The magnitude of an action potential is independent of the intensity of the stimulus and remains constant throughout its passage along an axon; what varies with stimulus intensity is the rate at which action potentials are produced.

What I wish to highlight in this process, as of special relevance to our object of study, is that the human retina, which as we have just seen has over 126 million photoreceptors, has only 1.0-1.5 million ganglion cells; and that whereas in the peripheral retina large numbers of photoreceptors are connected to each ganglion cell, which thus channels information on a relatively large part of the visual field, in the fovea there is approximately one ganglion cell for each cone (Figure 5). This privileged treatment afforded by the architecture of the visual perceptual system to information entering via the fovea helps explain why foveal vision, vision in the centre of the visual field, is so much sharper than the peripheral vision mediated by the retinal periphery.




Figure 4 Primary visual stream

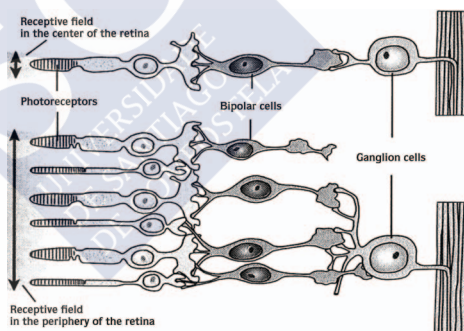


Figure 5 Ganglion cells in the fovea and the periphery of the retina

c) The primary and associative visual cortex

The vast majority of the information sent to brain by the retina is received in the striate cortex, or primary visual area (area 17 of Brodmann's map) of the cortex (Figure 6). In this region of the cerebral cortex the retina is mapped onto a mosaic comprising some 2,500 modules, each of which occupies a roughly 0.5 x 0.7 mm rectangle of cortical surface and contains approximately 150,000 neurons. Each module processes information on colour, luminosity, orientation, movement, spatial frequency, texture, retinal disparity and other properties corresponding to a single small area of the visual field. Together the cortical

modules process information from the whole of the visual field, but 25% of them are devoted to processing information from the fovea, i.e. information on what we are paying attention to and have centered our gaze on. The privileged treatment of foveally received information in the retina is thus echoed in the cerebral cortex.

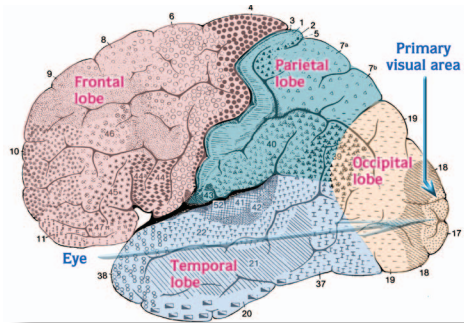


Figure 6 Primary visual area

It is indeed basically information derived from foveally received stimuli that leads to conscious visual perceptual identification of things, information from the periphery of the visual field being used to a much lesser extent. I can identify a clock without paying attention to any of its features in particular, but if I want to know the time I must fix my attention on the exact position of its hands relative to the numbers on its dial. Fixing my attention involves concentrating my gaze on that feature of the clock, so that the stimuli emanating therefrom be received by the cones of the fovea and consequently processed with the detail required by the desired perceptual identification. If I do not pay the clock this degree of attention, then the stimuli emanating from its hands and from the numbers on its dial may indeed be received by the photoreceptors of my retina, but since they will not be processed with the detail required by a representation of the time by the exact positions of the hands on the dial, I will not perceive what time it is, despite the stimuli being available for processing.

The primary visual cortex is essential for the conscious identification of objects and their properties, but is itself “blind”. Identification does not take place in the primary visual area, but in the visual associative areas. Upon arrival at the primary visual area, the information on the visual field is analysed by its modules, which then each send information on colour, luminosity, texture, spatial frequency, retinal disparity, orientation, movement and other properties of its portion of the visual field to the corresponding areas of the associative visual cortex, which are each specialized in processing one of these classes of information. It is at the end of this process, once the various types of information have been processed and integrated, that perceptual identification takes place. This associative processing is performed in parallel along two major pathways, the ventral stream and the dorsal stream (Figure 7). Signals travelling down the ventral stream, which terminates in the inferior temporal lobe, are integrated to extract information on the colours, shapes and identities of objects, while information on their motions and spatial locations is extracted in the dorsal

stream, which terminates in the post

Broadly speaking, the extraction of information from the signals sent by the retina is thus organized hierarchically, each region of the brain that is involved in the hierarchy processing information received from lower echelons and sending on the results of this processing to a higher echelon. However, as well as being passed on to the corresponding unimodal higher echelon of the visual associative cortex, these results are also sent to three other major cortical regions: the associative cortex of the frontal lobe, the paralimbic cortex (the gateway to the amygdala and the hippocampus), and the areas of multimodal convergence of the temporal and parietal lobes (Figure 8) (Fuster, 2003, p. 67). The connections with the frontal lobe supply information to the associative motor areas that control attention, planning, decision taking, etc.; the connections with the paralimbic cortex supply information to the areas involved in emotions and memory; and the connections with the areas of multimodal convergence supply information for intermodal association.

The fact that the processing of stimulus-induced signals involves the progressive ascent of a hierarchy of stages constitutes neurological grounds for thinking that the mind is modular in nature. But there are also numerous neurons with axons that transmit signals in the opposite, downward direction: from the three regions just mentioned to the associative visual cortex, from this latter to lower echelons in the hierarchy. The existence of these connections affords neurological support to the idea that prior knowledge must play an important role in the processing of stimuli. Because we know how to tell the time, we pay attention to the kind of stimuli that provide sensory information on the exact positions of the hands of the clock on its dial, the perceptual information we need to identify what time it is. If we did not know how to tell the time, that




Figure 7 Ventral and dorsal streams

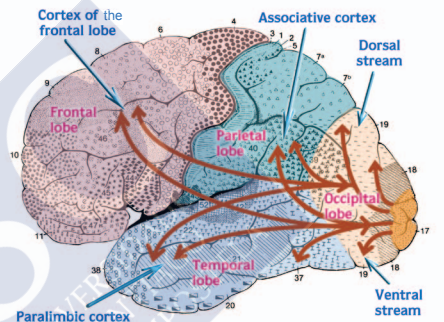


Figure 8 Cortical connections of the visual system

particular kind of sensory information would seem uninteresting, and we should consequently pay it no attention; and in not paying it attention, in not focusing our gaze to receive it, we should prevent its being processed with the detail required to identify the exact positions of the hands of the clock on its dial, and thus to identify what time it is.

From the foregoing brief sketch of the architecture of the human visual system we can draw at least three important conclusions.

a) What we perceive visually is not stimuli from the outside world (or from our own organism), but the result of their processing.

b) What we perceive visually does not depend only on the stimuli present in our visual field, but also on what kind of photoreceptors receive these stimuli. One kind of sensory information will be obtained if they are received by the cones of the fovea, and another quite different kind if they are received by the rods of the retinal periphery; and consequently, the perceptual identification of what is perceived will also be different.

c) What kind of photoreceptor receives the stimuli is in turn determined by the selective attention of the perceiver, which as we shall see in the following chapters is largely determined by his or her prior knowledge or categorial framework.





III. The categorial framework



Throughout this essay I shall maintain that the perception of adult individuals is categorial; in this chapter I shall examine the process of category formation (categorization), the constitution of categorial frameworks. Although the term “category” at once suggests the names of Aristotle and Kant, among others, I shall here defy philosophical custom by using it to refer to classes or kinds of things, events and properties, in the same way as it is used in cognitive psychology and anthropology. In the cognitive sciences categories are understood to be the result of organizing the stimuli received from the environment in meaningful groups by searching for common features. Thus we have the categories of dogs, cats, humans, chairs, colours, and other things, events, states and properties to be found in the world (Figure 9). Adapting this general characterization of categories to the use of the term “stimulus” established in Chapter I, we see that *categories are the result of organizing in meaningful classes, by searching for common features, the sensory information obtained by processing stimuli.*

No cognitive scientist doubts the existence of categories and processes of categorization, which are essential for the survival of the organism. Such processes are found not only in man, but also in the immense majority of animals. A dog has no difficulty in distinguishing other dogs from cats or birds, food from what is not food, and so on. Like us, dogs organize the sensory information derived from the processing of stimuli in meaningful classes, that is, classes that are meaningful for them, and which in many cases are quite similar to ours.



Figure 9 Graphic examples of categories

The behaviour of a dog towards another animal depends on whether that animal is another dog, a cat, or a bird; and it behaves similarly towards things of the same type or category. To judge by how they interact with their environment, most animals, besides identifying types of things, can also identify individuals; but like humans, they identify an individual as being a certain kind of thing, a member of a certain category (which is not to say that they are necessarily

explicitly aware of doing so).

In the cognitive sciences the problem is not whether categories exist or not. As I stated above, no-one doubts that they do. The problems posed by categories in the cognitive sciences are of a different kind, and have to do with, among other things, their origin, their constitution, their universality and, especially, the mechanisms by which they are identified and represented by organisms. Are there *a priori* categories, or are they all *a posteriori*, constituted in the course of development? Are there *basic* categories, as postulated by Eleanor Rosch? And if it is indeed possible to distinguish between different levels of category (basic, subordinate and superordinate), at what age do children acquire the categories of each level? And as regards universality, are there categories that are found in all cultures (for example the primary colours red, yellow and blue), or are categories specific to each culture? And so on.

Perhaps the most important problems of the cognitive sciences, and those to which most attention has been paid, are those concerning the mechanisms by means of which categories are identified and represented. Are they represented by sets of features, or by necessary and sufficient properties? By prototypes? By a more or less vague memory of the examples that have been perceived? Atomically, or in conceptual networks? All these issues concerning categories are problems that confront cognitive psychology and anthropology, and philosophy too; but as I noted above, what nobody questions is their existence, the fact that human beings, in the course of their development, learn to organize sensory information in classes or types of thing, event, or property, etc.

Categorization simplifies the perceived environment and reduces the burden on memory. It means storing and retrieving information in a way that transcends singular entities, affording access to more information for less cognitive effort. Through categorization, we cognitive beings taxonomize the world, learning to organize sensory information in classes or types of things, and this allows us to apply to each member of the same class or category the information pertinent to that class, a capacity that is fundamental for our survival. When we visit a city we have never been in before, we drive through its streets and, using a map or whatever, we arrive at the hotel we booked a room in. We can do this because, even though everything around us is new to us, the sensory information derived from the processing of the stimuli received by our sensory receptors is experienced in the process of perception, integrated in the configuration of our pre-existing categories, the categories of pavements, streets, traffic lights, buildings, cars, men, women, etc. It is in this way that the knowledge that is available in relation to each of these categories or types of thing guides our actions.

If we only knew of and about singular entities (people, chairs, tables and trees that we already know), then on arrival at the new city we would be in the same

situation as a newborn baby. Faced with a tree on the pavement, for example, we would not know whether it was a living being or an artifact, or whether it would bite us, burn us or shriek if we approached it. But on identifying it as a tree, as a member of the category of trees, our knowledge about trees will guide our behaviour. In particular, we will know that the tree is not going to bite us, burn us or shriek, and we will also know that it is not going to step aside to make way for us either, so if our purpose is to get back to our hotel it is we who must avoid walking into the tree.

Through categorization we semiotize the world, we endow it with meaning – on the understanding that in the use I wish to make of the term “category” here, the categorial nature of objects, events, properties, etc., is not something independent of them, but is on the contrary their way of being objects, events and properties for cognitive beings. It is thus invested in their categorial clothing that the entities of the world are identified in the processes of perception by cognitive beings who possess the corresponding categorial systems.

The sensory information leading to a categorial identification can be very varied. I can identify the animal in the garden as a bird seeing it perch on a twig, open its wings to fly, or hop on the lawn; or hearing it sing; or seeing it from a position different from that which I now occupy. In each case the sensory information leading to its perceptual identification as a bird is different; but in all of these cases the same information could have led either to its identification simply as a bird or to its identification as a certain kind of bird. Furthermore, in many circumstances the sensory information leading to the categorial identification of something as a certain kind of thing can be considerably less precise than the categorial identification itself. Suppose I am looking at cars travelling along a road a mile or two away. The shapes, colours and sounds (if I hear them) that lead me to identify each passing vehicle as a car are quite indefinite – I could say very little, if anything, about this sensory information – yet I have no difficulty in identifying each of those passing vehicles as a car; it would be no problem for me to count how many pass in the next quarter of an hour, say.

It is not only in cases like the one I have just described that the sensory information leading to the categorial identification of perceived items is hazy; it occurs in the immense majority of perceptions. In our everyday interaction with our environment it is categorial perceptual information that is relevant, and the corresponding categorial identification can in each case be based not only on different sensory information, but on sensory input to which we pay no great attention. We will subject our sensory information to the processing that is necessary for the corresponding categorial identification, but to no more than is necessary. If I propose to cross the street, the exact colour of the post supporting the traffic lights is of no consequence, nor its height, or design, or

any property other than that the light showing is red, green or amber; so all those other properties, which are irrelevant for the type of action I propose to perform, will not be processed with the detail required for their categorial identification. Similarly, if what I want is to open a door or drink a glass of wine, all the stimuli in my visual field that are irrelevant for opening the door or drinking the wine will be processed with less than that degree of detail that is necessary for the corresponding categorial identification, and the sensory information derived from their processing will be more or less hazy.

There are, of course, circumstances in which what is of interest is the sensory information itself, as in a philosophical discussion between conceptualists and phenomenals about the content of our perceptual experience. The phenomenalist will draw attention to the phenomenal content being the richer, in that our perceptual experiences are much finer-grained than the categorial content that is expressible in concepts. And so they are, in those cases in which what we pay attention to is precisely the sensory information. The problem for both phenomenals and conceptualists stems from their striving to reduce the content of our perceptual experiences to just one of their two aspects, and failing to realize that the processes of perception involve both kinds of information, sensorial and categorial, as we shall see in Chapters IV and V.

What I am trying to show with all this is that it is the categorial content of our perceptual experiences that is truly important in the processes of perception, even though it does not exist without the sensorial information with which it is perceptually associated. If an eskimo or a skier can distinguish among different kinds of snow it is because there are sensory features that permit those distinctions; but those features, on which the eskimo and the skier base their different types or categories of snow, can be totally unperceived by someone who has never bothered about those distinctions, and who consequently lacks the corresponding categories. Since those sensory features are of no interest to him, he will not pay them the necessary attention, and their processing will consequently not be so detailed as to lead to identification.

In the same way that, for any given individual, the sensory information leading to the categorial identification of a certain item can and does vary from one perception to another, without the identification of the item necessarily being affected, so too, different individuals, to the extent that they share the “same” categorial frameworks, can make the same categorial identifications of the same items, even though they make those categorial identifications on the basis of different sensory information. At the level of uncategorized sensory information we live in private worlds; but not at the level of categorization. Where I see a bird in the garden, so too my neighbours, if they look at it, will also see a bird, even though the sensory information leading to the identification of what we each see as a bird is for each of us not quite the same as for the

others. And this shared categorial framework allows interaction among us; it is those “same” categorial contents that, through their representation in language, allows us to communicate with each other.

Now that we have come to the subject of “representation”, it is time for a further terminological clarification that I believe to be of great importance for the analysis of perception. Although it is quite usual in the cognitive sciences to use the terms “category” and “concept” as equivalent, I am going to distinguish between these closely related notions. Following Gregory Murphy (2004) and Jesse Prinz (2004),¹ I shall use the term “category”, as I have in fact been doing, to refer to “how things are” in the world, to their nature as things, and hence not for anything that may exist without them. The category of chairs is constituted by all chairs, by each and every one of them; the categories of trees, dogs and men by all trees, all dogs and all men, respectively; and so on. I shall use the term “concept” to refer to a mental representation of a category. Thus categories and concepts, though closely linked are distinct: categories are denoted by the concepts that are most fundamental and closest to empirical experience.

As we shall see in greater detail in Chapter VI, in the processes of perception things present themselves to us, we have direct access to them as they exist in the world processed by the architecture of our perceptual system. Accordingly, it is erroneous to think of the categorial content of our perceptual experiences as *representational*; this content is *presentational*. What is representational, and constitutes our concepts, is the more or less vague memories of this categorial content contained in the sensory information.

I remarked above that no-one in cognitive science doubts the existence of categories. On the other hand, one of the thorniest and most serious problems of cognitive science concerns their representation.

According to the “definitional theory” that was more or less current until the 1970s, concepts are fixed by definitions. Piaget’s experimental research on the cognitive development of children was carried out taking for granted that concepts are mental representations for which precise definitions are available. The intensional content of a concept is in this theory determined by the set of properties that are common to all the members of the class or category to which the concept refers, together with the set of differences that differentiate the members of that class or category from those of others.

From the standpoint of logic, to fix or characterize concepts by means of definitions has many advantages. In the first place, if correct, it furnishes us

¹ ‘In general, I try to use the word *concepts* to talk about mental representations of classes of things, and *categories* to talk about the classes themselves’ (p. 5). Murphy, Gregory L. (2004): *The big book of concepts*. Cambridge, Massachusetts, MIT Press.

‘The concepts are mental representations of categories that are or *can be* activated in working memory’ (p. 149). Prinz, Jesse J. (2004): *Furnishing the Mind. Concepts and Their Perceptual Basis*, Cambridge, Massachusetts, MIT Press.

with the means for precise identification of the members of the corresponding category. Given the definition of a dog, one can examine the parts of the definition to determine whether something is a dog or not, since all and only those entities that possess all the attributes specified in the definition will belong to the category of dogs. Possession of all these attributes is a necessary and sufficient condition for belonging to this category: if something possesses all the attributes required by the definition of dogs, it will *be* a dog, and if it lacks any of these attributes it will not be. The definitional characterization of concepts, if valid, thus allows clear segregation of the things that belong to the corresponding category from those that do not.

Secondly, a definitional characterization of concepts allows one to establish precise relationships among them, to draw conclusions, to justify certain kinds of analyticity, and so on. If all the attributes employed to characterize the concept of animal are also found in the definitional characterizations of the concepts of dogs, cats and birds, I may conclude that dogs, cats and birds all belong to the category of animals. Similarly, if the concept of extent figures among the attributes employed to characterize the concept of body, I may conclude that the statement “all bodies are extended” is an analytic truth.

According to the definitional theory of concepts, their empirical referents, categories, exhibit three fundamental features.

a) They are characterized by a set of jointly necessary and sufficient conditions, properties or attributes satisfied or possessed by all their members, so an individual acquires a category when he has a mental representation constituted by all these conditions, properties or attributes.

b) All the members of a category are equally representative of it, since anything fitting the definition is as much a member of the category as any other member.

c) Categories are discrete, there is nothing intermediate between two categories that is not itself a category. It is always theoretically possible to determine whether something does or does not comply with a categorial definition, even though we may not know how to carry out that determination; if it does comply, it is a member of the category, and if it does not, it is not a member.

This view of concepts received severe criticism during the second half of the twentieth century, on both theoretical and empirical grounds. In regard to theoretical considerations, the “second” Wittgenstein was already questioning

the possibility of providing definitions for most concepts; and though his arguments are basically negative, anyone who tries to contrive a definition of some of our commonest concepts – dog, cat, tree, chair, table – will soon find the same or similar difficulties as were brought out by Wittgenstein's analyses in the case of the concept of game. This is not to say that this traditional idea of concepts is totally obsolete and cannot be used to characterize certain kinds of concept, such as formal concepts and a limited number of scientific concepts; but in general the groups of concept that we can characterize by definition are the exception rather than the rule.

On the empirical side, it was the pioneering anthropological and psycholinguistic research of Brent Berlin and Eleanor Rosch on categorization processes (Berlin *et al.*, 1969; Rosch *et al.*, 1975) that most contributed to the downfall of the definitional characterization of concepts in the cognitive sciences. Contradicting the three features presented above as characteristic of the definitional view of concepts, empirical research shows that in general, when a group of individuals are asked to list the attributes something must possess in order to belong to a certain category, those listed by different individuals only agree in part, and the lists never, or hardly ever, manage to provide necessary and sufficient conditions for membership in the category; and if subjects are instead asked to classify a series of cases as belonging or not belonging to a certain category, they will all agree in some cases but not in others. They will all agree, for example, that apples, oranges, pears and peaches are fruit, but may disagree as to whether melons, watermelons, dates or olives are; and similarly in regard to the members of many other categories (Rosch and Mervis, 1975). These findings show that not all the members of a category are equally representative of it; some are more representative or prototypical, and others less so, a fact that the definitional characterization overlooks or cannot account for. If we ask a group of individuals to give us some examples of fish, it may well be that they all mention trout, sardines or bream, and very few of them eels, lamprey or sharks; it is as if these latter were not such good examples of the category of fish as are trout, sardines or bream. And when experiments measure the time taken by subjects to identify things, it is found that they recognize the more prototypical members of a category much faster than the less prototypical or more peripheral. We identify a cow or a sheep as mammals faster than a dolphin or a whale. Moreover, in developmental studies it is found that children learn to identify the members of a given category faster if the examples used to teach them are prototypical than if they are atypical.

Finally, there are things that possess some of the typical characteristics of two or more categories without possessing all the typical characteristics of any of them. Is a seat with three legs and a low back a stool or a chair? Are tomatoes and olives fruit or vegetables? Empirical research shows that the answers to

these and similar questions vary widely from one person to another, and even that the same person can give different answers at different times and not because they have not paid attention to the question, but because on one occasion they focus attention on one aspect of an object and on another occasion on a different aspect. Tomatoes seem to be fruit as regards taste and form, but if we think of the way they are seasoned they are like vegetables.

In short, empirical studies show that categories possess properties that are quite different from those predicted by the definitional theory of concepts.

a') Concepts are constituted by a rather ill-defined set of properties or characteristics with different weights in the determination of the corresponding category. Our minds appear not to be so analytical as the definitional theory requires; as Wittgenstein pointed out, what allows us to identify the members of a category may be only a subtle “family resemblance” shared by all its members.

b') Not all the members of a category are equally representative of it. Some are more central or prototypical than others. It is closer to the use we make of concepts to think of the members of a category as arranged in concentric circles with different degrees of centrality or membership in it, rather than as all belonging to a single uniform set that is equally well represented by them all.

c') The boundaries between categories are often fuzzy.

These three features that appear to be shared by categories are difficult to reconcile with the definitional theory of concepts, and have led to the emergence of alternative theories. These alternatives seem to be more in keeping with the results of empirical research, but fall short of being entirely satisfactory. Nobody doubts that categorization occurs, or that we access a categorized world, but no theory has yet emerged that satisfactorily explains the mechanisms by which the human mind – or, if you will, the architecture of our perceptual system – builds our mental representations of categories and identifies their members in the processes of perception. It may be that a satisfactory answer to this kind of question can only be attained using more information about the neurological basis than is currently available.

In the remainder of this chapter I am going to present, albeit very succinctly, the three most important theories that have emerged in the cognitive sciences as alternatives to the definitional theory of concepts. These are *prototype theory*, *exemplar theory*, and the *theory of theories*.

Prototype theory

Eleanor Rosch was not only one of the earliest critics of the definitional theory of concepts, supplying abundant empirical evidence against it, but also laid the foundations for the construction of an alternative, the *prototype view*. Many commentators on her work have interpreted the prototype view as the suggestion that every category is represented by a kind of prototype or ideal specimen constituted by the set of attributes or properties that are generally found to be possessed by the members of that category. Thus the category of birds would be represented by an ideal bird that would have the following attributes, among others: animal, medium size, covered in feathers, with two legs, two wings, two eyes and a beak, capable of walking, singing and flying, etc. So interpreted, the theory of prototypes adequately accounts for the phenomenon of unequal representativity. In the real world, not all members of a category resemble the prototype to the same degree; some are more like it, some less, which explains why, in the process of perception, some are experienced as more central to the category and others as less central. The most typical specimens are those that most closely resemble the prototype, and the less typical those that least resemble it. Herons and cormorants conform less well to the ideal prototypical bird than do sparrows or robins, and there are also birds about which we might have serious difficulty in deciding whether to include them in this category or not. To this extent, the theory of prototypes appears to offer a better account of categories than did the definitional theory, according to which all the members of a category are equally representative of it.

Nevertheless, this interpretation of the theory of prototypes, according to which categories are represented by an ideal prototypical specimen, poses problems. In the real world there are categories in which the variation among the members is minimal compared to the variation among members of other categories. Cats, for example, are less varied than dogs or birds: the variation among different kinds of bird or different kinds of dog is much greater than the variation among cats. These between-category differences in the extent of within-category variation are not explained by the theory of prototypes, which postulates that each category is represented by a single ideal prototype.

An alternative interpretation of Eleanor Rosch' theory of prototypes has it that categories are represented by a summary description of the category as a whole, rather than by a single ideal prototype (Hampton, 1979; Smith and Medin, 1981). The problem with this approach is its vagueness. What should we understand by "summary description"? To resort to Wittgenstein's "family resemblance" to describe what is shared by the various members of a category, as do Rosch and Mervis (1975), hardly removes this vagueness. The most plausible interpretation consonant with this vagueness is that categories are

represented by a description comprising most of the attributes of members of the category, and in which not all attributes have the same weight. Thus the attributes of having wings and feathers have greater weight in the representation of the category of birds than the attributes of capacity to fly, or of being carnivorous or herbivorous, but without the possession of any group of these attributes constituting a necessary and sufficient condition for identification of the category. The attributes found in a majority of the members of a category, though not in others, will have more weight in the representation of the category than others that are less common or are shared by members of other categories. But what about the representation of attributes that are contradictory, or exhibit great variation, yet are found in different members of a single category? In the category of dogs, for example, we find dogs of very different sizes and colours, some with long hair and some with short or practically no hair, and so on in regard to numerous other attributes. These are all problems for which the theory of prototypes appears to have no explanation, though the interpretation of this theory in terms of a summary description of attributes that have different weights in the representations of categories seems to be more on target than the interpretation in terms of the construction of a single ideal prototype, the ideal prototype being unable to account for there being greater variation among the members of some categories, such as trees, birds or dogs, than among others that exhibit great uniformity, such as cats, tigers or lions (in spite of which, the supporters of the theory of prototypes generally associate the notion of prototype with a single ideal specimen, not with a list of properties or attributes).

There is a rather more elaborate version of the theory of prototypes according to which categories are not represented by unstructured lists of individual properties with different weights, but by a structured set of properties (Cohen and Murphy, 1984). It is not as a mere collection of individual attributes that properties we observe in dogs (e.g. The possession of one head, four legs, two eyes, two ears, a tail, etc.) allow us to identify them as dogs, but by virtue of their being mutually related in a certain structure. A dog is not an amorphous heap of properties, and cannot be represented by one. As well as a set of properties with different weights, one needs something else for this set of properties to be able to be representative of a category, and that something else is the structure in which the properties are mutually related. A dog with its head, eyes, ears or tail located in unusual parts of its body would at least raise eyebrows, and we might even doubt whether it was really a dog; and this implies that the structural organization of properties must also form part of the representation of a category.

As the foregoing paragraphs illustrate, there are various different ways of interpreting the theory of prototypes, none of which is totally satisfactory. Depending on the type of experiment being carried out, it may be sufficient to think of the representational content of a category as an ideal prototypical

specimen, or it may be more useful to think of it as a set of properties with different weights in the representation. And there will be occasions on which not even the set of properties together with the structure they “inhabit” will seem adequate. In conclusion, no available interpretation of the theory of prototypes affords a totally satisfactory explanation of the way our brain manages to represent categories and use that representation to identify their members so efficiently.

Exemplar theory

Exemplar theory, first put forward by Medin and Schaffer (1978), provides an alternative to the points of view represented by the definitional theory and prototype theory. For the exemplar theory of concepts, mental representations of categories are not constituted either by definitions that cover all and only the members of the category in question, or by sets of properties that are exemplified to a greater or lesser degree in each of the members of the category. This theory in fact rejects the idea of a unitary representation of conceptual content; according to exemplar theory, a person’s concept of dog is nothing but the more or less vague memory of all the dogs that this person has perceived in the past. In a sense, this point of view implies the negation of concepts, since it denies the existence of a unique representation valid for all the various members of a category. An individual faced with a new specimen consults his or her memory to determine which of the known exemplars of the various categories is most similar to the new one, and then proceeds to ascribe this new one to the corresponding category. We identify the dogs, cats and birds that we perceive as belonging to their respective categories because of the similarity of each with the members of that category that are already stored in our memories. According to this approach, what we now perceive as a duck is identified as such because, among the various kinds of bird we have perceived in the past, it most resembles the exemplars that are grouped together in the category of ducks.

The exemplar theory of the conceptual representation of categories, unlike the definitional theory, accounts for the phenomenon of variable centrality, for the fuzzy nature of many categories, and for many other features of category representations that have been revealed by empirical research. An exemplar of a category will be more prototypical or central for that category, and thus easier to identify, the more closely it resembles the exemplars of the category that are already known. If I live in an area in which sparrows are one of the most abundant species of bird, I will identify a robin as a bird more easily than an eagle, a goshawk or a barn owl, because robins resemble sparrows much more

closely than do these other birds.

The exemplar theory also accounts satisfactorily for fuzziness in the identification of certain kinds of exemplar. Although whales and dolphins resemble fish because of their morphology and the medium in which they live, consideration of how they reproduce and feed their young shows their resemblance to other mammals. Whether they are identified as members of one or the other category will depend on which aspects are dominant in their perceiver's thoughts.

There are a number of other arguments that appear to support exemplar theory. For example, it is logical to suppose that the first time an individual perceives a member of a certain category, the first time he or she thus makes direct perceptual contact with that category, it will be the memory of that first exemplar that is stored as representative of that category - unless we accept that humans have an innate capacity equivalent to Aristotelian abstractive induction (*Posterior Analytics* 71a-8), making it possible to apprehend the universal implicit in the particular. Again, representations stored in our memory are, at least in some cases, representations of particular exemplars, as postulated by the theory; this is the case of the things with which we are most familiar - our home, our parents, brothers, sisters and friends, and a multitude of other things that we are all conscious of preserving a particular representation of.

Yet in spite of its agreement with the results of experimental research, and its ability to explain cases like those mentioned above, a number of critics maintain that exemplar theory contradicts our most immediate conscious experience concerning general concepts. In the perceptual process of identifying something as a cat, a dog or a tree, we are not conscious of remembering particular examples of cat, dog or tree that we have perceived in the past; rather, we conjure up, as it were, a general idea of those different categories or types of thing. Similarly, when we make use of a concept in thinking, what appears to be active in our memory is a unitary representation of its category, not representations of each individual known exemplar.

It may of course be argued that this general knowledge of categories may be just what is common to the exemplars we remember, and that in any case we should not forget that both the use of concepts in perception and their use in thinking are automatic processes that are partly pre-conscious and extremely fast, and that our conscious experience may accordingly not be a reliable touchstone for their analysis. But there is a further difficulty to be faced by exemplar theory. In the constitution of the mental representation of a category, what is it that counts as an exemplar? Is it each of our perceptual experiences of a member of the category, regardless of whether they correspond to the same or to different members? Or is it only experiences of different individuals that are stored in memory? There is no simple answer to these questions, given the

intrinsic discriminatory limitations of our perceptual system and memory. How can I know, on the basis of a simple perceptual experience, that the poodle I saw in my neighbour's garden at five o'clock this afternoon is the one I have just seen, an hour later? If I had seen poodles with significantly different sensory features (different in colour and size, say), then I could confidently assert that they were different poodles; but if the differences between the two sightings were small, I should probably not have noticed them as differences, or have forgotten them if I noticed them, and in consequence I should not know whether it was the same poodle or two different poodles. More difficult still to answer: is it possible to decide, on the basis of a simple perceptual experience, whether the robin that has just alighted on the window sill of my study is the same one that alighted there yesterday?

Finally, there is a difficulty facing both exemplar theory and prototype theory, which both appeal to the problematic notion of similarity. The identification or recognition of new members of a category is in both these theories explained by their similarity to representations stored in memory: a prototype in prototype theory, exemplars in exemplar theory. But whether x is deemed similar to y will depend on the attributes of x and y that are deemed relevant, and on their relative degrees of relevance. Big Ben belongs to the category of clocks if what weighs most in our decision is knowing the time, but we will identify it as a tower if greatest weight is given to the edifice. Medin and Schaffer (1978) tried to provide a measure of similarity in their "Context theory of classification learning", but without much success.

Theory theory

Just like prototype theory and exemplar theory, theory theory emerged in relation to experimental research. But whereas prototype theory and exemplar theory were prompted by the failure of classical definitional theory to explain empirical results, theory theory emerged as a kind of reaction to prototype theory and exemplar theory, which both treat categories and their mental representations (concepts) as if they were monads, overlooking the influence of available knowledge in the processes of categorization. Theory theory may be regarded as complementing prototype theory and exemplar theory, rather than as an alternative to them.

For theory theory, concepts are part of our general knowledge of the world, and they are not acquired independently, in isolation, as is assumed in most experimental designs, but as parts of a conceptual network integrating all our knowledge of the world (Murphy, 2004). When someone acquires knowledge

of a new category, such as a new breed of dog, that knowledge does not emerge from nowhere. That person's existing knowledge about dogs, animals, and living beings in general, and about the various ways they interact with their environment, influences the new category; and conversely, the acquisition of a new category involves modification of that person's conceptual system.

The knowledge someone has about how human beings and other animal species feed and care for their young will allow him or her, upon first seeing a hen with her chicks, to identify the chicks as her young, something that person might not be able to do in the absence of that knowledge, given the scant similarity between chicks and adult cocks and hens. Contrariwise, the acquisition of new information about the environment (for example, that whales and dolphins are mammals, or that earthworms and snails are hermaphrodites) will involve reorganization of the biological knowledge of the person acquiring this information. From this point of view, part of the processes of categorization may depend both on stimulation and on the prior knowledge at the perceiving subject's disposal. As we saw in the example of the hen and her chicks, the identification of these latter as members of the category of hens depends both on sensory stimulation and on inferential processes based on available prior knowledge.

The same animal that I might identify in the wild as a wolf, I might identify as a dog if I see it walked in an urban park. My knowledge of wolves and dogs, and of the different environments they inhabit, would influence my categorial identification of the animal. If I see something I identify as a bird of prey sitting on a branch of a tree, and it suddenly happens to fall to the ground with a metallic "boing", my prior knowledge about birds will allow me to infer that the animal I identified as a bird is no such thing, but a mechanical device. In short, categorial identifications are the result of the processing of information in both directions, from bottom up and from top down. Stimuli from the outside world or from our own organism provide us with sensory information that is organized in the light of the representations previously stored in memory, thereby giving rise to the categorial identification of the perceived items. Influenced by our prior knowledge of the configuration of faces seen right way up, which is how we generally see them, we identify facial parts (eyes, mouth, nose, ears, cheeks) much better when we see a face right way up than when we see it upside down. The information stored in memory and the information received from the outside world can interact in this way because of the two-way organization of our neuronal connections. In fact, the primary visual cortex sends more projections to the lateral geniculate nucleus of the thalamus than it receives therefrom, and receives more from the associative visual areas than it sends thereto. These reciprocal connections allow reiterative processing of information, with information being exchanged in both directions, thereby improving the

information obtained from stimulation and prolonging its residence. Because I know there are different types of snow, I pay attention to those features of snow that allow these types to be distinguished; and because I pay attention to these characteristic features, I see the snow as conforming to one or another type. If, on the other hand, my life has been led in an environment in which snow is a rarity and it is pointless to distinguish among different types or categories, my perception will limit itself to consideration of the kind of sensory information that is sufficient for seeing its source as conforming to the basic category of snow. And what goes for snow, goes for everything else in our environment.

Theory theory emphasizes the role of prior knowledge in both the acquisition of new categories and the perceptual identification of their members. We pay attention to those features of the world that prior knowledge tells us are relevant; and because we pay attention to these features, we make certain categorial identifications and not others. This does not mean that categorization or categorial identification processes are arbitrary. If there were no sensory features differentiating different types or categories of thing, those categories could not be established; but it may well be that many sensory features that might be taken into account never actually attract our attention, and we consequently fail to see them as conforming categorial identifications that might otherwise be possible. A shepherd pays attention to sensory features that allow him to discriminate among different kinds of sheep, and because he pays attention to those features he distinguishes one from another. The layman in this categorial field simply sees sheep which are all different from each other, but not so different as for him to cease identifying them all simply as sheep, as specimens of a single category.

Theory theory is possibly correct in stressing the influence of prior knowledge in the acquisition of new categories and the identification of their members, and in maintaining that concepts are not monadic representations – as prototype theory and exemplar theory seem to imply – but part of networks of connections in which they act, as it were, as nodes. It is through these nodes that the conceptual networks make contact with their referents, but the content of the nodes is not determined only by their categorial denotanda in the world, but also by their integration in the whole constituted by the network.

It is clear that theory theory is right to draw attention to an aspect of our knowledge that both prototype theory and exemplar theory overlook. But it is no less clear that theory theory is unsatisfactory in regard to the acquisition of the first mental representations. Not all concepts can be generated with the aid of prior knowledge.

As has been pointed out many times, it is a fact that no-one questions, because it is so obvious, that in the processes of perception human beings organize the sensory information derived from stimulus processing in classes or types of thing, property, event, etc. The problem, as we have seen in the foregoing, is to explain how our perceptual system creates and stores representations of these categorial contents, and

how, once stored, they are used to identify the categories to which new specimens belong.

It may be that the answer lies in recognition that the three theories sketched above are mutually complementary, and that for mental representation of categories the architecture of our perceptual system very probably makes use of one or another mechanism depending on the particular case in hand and the concurrent circumstances. In fact, as noted above, we have mental representations of particular things, but we also have representations concerning knowledge in general; these latter seem more like prototypes or unitary schemes than memories of perceived specimens, yet it is certainly true that available knowledge influences both the acquisition of new categories and the perceptual identification of their members.

It may be that a satisfactory solution to the problem of the mental representation of categories must await further understanding of the neurological basis of cognition. At present, all that has been achieved by prototype theory, exemplar theory and theory theory has been to provide explanations compatible with two of the phenomena that all empirical studies about categories appear to agree on: *a*) the existence of different degrees of typicalness (not all the members of a category are equally representative of it); and *b*) the existence of fuzziness (there are specimens on the borderline between different categories).

The taxonomic organization of categories

A group of equally important problems, that have likewise been the object of extensive experimental research, derive from the taxonomic organization of categories, a kind of nested structure with levels corresponding to degree of abstraction, each nested category being a subset of the higher-level category to which it belongs. Thus the category of animals includes the categories of cats, dogs and birds, and the category of dogs in turn includes the categories of poodles, chihuahuas and pekineses. This taxonomic structure organizes categories in multiple levels: in the example just given, the category of animals is a subset of the broader category of living beings, which is in turn a subset of entities in general; and many intermediate categories can be established. In spite of this, since Eleanor Rosch introduced the expression “basic level” (Rosch *et al.*, 1976) to refer to categories of intermediate level – in our example, the categories of dogs, cats and birds – studies concerning the taxonomic organization of categories generally distinguish just three levels of categorization: basic, subordinate, and superordinate.

That Rosch regarded categories of intermediate level as basic was due to this type of category having considerably greater cognitive importance than those of the other two levels, by which I mean that they exhibit features such as the following.

a) Objects are named more quickly when reported at this level than at the other levels (Rosch *et al.*, 1976; Murphy and Smith, 1982).

b) Objects are preferentially designated using their basic-level names (Rosch *et al.*, 1976; Berlin, 1992).

c) The number of characteristics listed by experimental subjects at the basic level is much greater than at the superordinate level, and only slightly less than at the subordinate level (Rosch *et al.*, 1976).

d) In the course of development, the names of basic-level categories are the first to be learnt, at around age 2 or 3 years, whereas those of superordinate categories are learnt at around age 4 years and those of subordinate categories at around age 5 years (Markman, 1989).

e) Basic-level categories are the most universal, in the sense of there being names for them in practically all cultures (Malt, 1995).

f) In general, speakers of all languages possess more words for basic categories than for subordinate or superordinate categories, even though the number of subordinate categories is logically greater (Berlin, Breedlove and Raven, 1973).

g) There are more one-word names for basic categories than for subordinate categories (Berlin, Breedlove and Raven, 1973).

h) In all languages, the names of basic categories tend to be shorter than those of others. As the frequency of use of a word increases, it tends to become shorter (Zipf's law in linguistics)

All these features reveal the cognitive preponderance of intermediate-level categories over those of the subordinate and superordinate levels; and the most widely supported explanation, though far from definitive, points to the fact that basic categories are highly informative and distinctive. Although subordinate categories are more informative, they are much less distinctive. To discriminate among subordinate categories it is necessary to pay attention to subtle characteristics that are far less striking than the morphological differences among intermediate-level categories. The members of an intermediate-level category usually share a common morphology that is markedly distinct from those of the members of other intermediate-level categories. The morphology of a dog, for example, differs markedly from that of a cat, a cow or a bird, while all

the members of each of these categories bear a striking “family resemblance” to each other.

In spite of all this, many cognitive scientists are still reluctant to consider intermediate-level categories as “basic”, since the categories of this level are not always dominant. For persons living in an urban environment, with a lifestyle quite far removed from nature (as is the case of 21st century Western societies), the dominant categories of plants and birds are no longer those of intermediate level but those of a higher level (Mandler and McDonough, 2000); while experts in a given field can process lower-level categories relevant to their field of expertise as efficiently as those of intermediate level (Tanaka and Curran, 2001). These findings somewhat detract from the “basic” nature of intermediate-level categories.


The ontological status of categories and concepts

If the term “category” is understood, as here, as referring to a type of thing obtained via stimulus processing by organizing the sensory information derived therefrom in meaningful classes through the extraction of shared features, then categories will be understood as having an ontological status quite different from that which corresponds to an understanding of categories as also embracing their mental representations (concepts), as is usual in the cognitive sciences. As I have been using the term here, a category is not something that exists independently of the particular things, properties or events that it encompasses. A category is the way in which things are things, properties are properties and events are events for cognitive beings. As we have seen in the foregoing, their constitution involves not only stimulation from the outside world, but also the architecture of the perceptual system of cognitive beings, and the interaction of these cognitive beings with their environment. In the formation of the categories of dog, bird and tree, sensory information derived from the processing of stimuli may be what is important, but if we consider the categories of vehicle or food, for example, what is important is their function, and for ethical or social categories the important determinants are a system of values and a social organization. Thus the way in which categories are things, properties, events, and so on, in some cases depends more on sensory information and in others on the functions of the members of the category, always in accordance with the set of factors involved in the interaction of cognitive beings with their environment. Different subjects, in different environments, may thus establish different systems of categorization, although some kinds of category, such as object or property, are found in all cultures.

These considerations entail what in my opinion is the ontological status of categories. Categories are constituted in the interaction of cognitive beings with their environment, and are the result of the processes of taxonomization pursued by these cognitive beings in this environment. This is not to say that they are subjective or lack empirical underpinnings. If cognitive beings have organized sensory features as they have done, it is because this categorization of the world has efficiently favoured their survival. In the process of development, a system of categorization that is inefficient, or is less efficient than an alternative system, is modified. This is what children do in the first years of their development, and what we adults continue to do whenever we acquire new data that subvert our pre-established system of categories.

Thus categories, in the sense given here to this term, are not things that exist independently of things themselves and the way they *are* things for cognitive beings. The situation is very different for the ontological status of concepts, the mental representations of categories. The ontological status we afford to concepts will depend on what we think about how the architecture of the human cognitive system constructs and stores categorial information; in other words, on our point of view in regard to representational contents. To think of the representational content of a category as constituting a unitary representation, in accordance with the definitional theory or prototype theory, is not the same as to think of it as a collection of stored memories of examples, in accordance with exemplar theory. And to consider those representational contents as self-sufficient monads is not the same as to consider them as forming a constituent part of a representational network. Consequently, all we can say in regard to the ontological status of mental representations of categories is that at present we have no satisfactory account of the mechanisms by which we cognitive beings perform the constitution, representation and recognition of the categorial beings that the things, events and properties of the world present themselves as in the processes of perception. And in the absence of a satisfactory account, we must postpone assignment of any particular ontological status to the mental representations of categories.

That said, it must be pointed out that for my purposes here, the ontological status of concepts is somewhat irrelevant. My aim is not to give an account of the mental representations of categories, but to provide a satisfactory explanation of perception and its contents. For this purpose, the fundamentally important fact is something that no cognitive psychologist or anthropologist has ever questioned: the human capacity to learn to recognize similar items and group them in *categories*, a task at which the human brain has shown itself to be especially expert.



IV. Subcategorial and categorial perceptual information



In Chapter 2 of *Being No One. The Self-Model Theory of Subjectivity*, Thomas Metzinger exploits an analysis of Diana Raffman's² to introduce the distinction between categorial and subcategorial qualia. This distinction is important, although it is perhaps not appropriate to use the term "qualia" to refer to both categorial and subcategorial information, since, as we shall see in this chapter, only subcategorial information has the subjective character that is generally attributed to *qualia* by their proponents. In my view, the term "qualia" could only meaningfully be applied to subcategorial information, not to categorial information. Consequently, while accepting Metzinger's distinction between two types of information, in what follows I shall use slightly different terminology: instead of the term "qualia" I shall speak of "information", distinguishing, *à la Metzinger*, between categorial and subcategorial information. Categorial information is constituted by attentionally, behaviourally and cognitively available content, whereas subcategorial information is constituted by content that is available only attentionally, and behaviourally, not cognitively.

To clarify the differences between these two types of content, I shall begin by briefly characterizing subcategorial information, and then proceed to describe how, in the processes of perception of adult subjects who already possess a categorial system, categorial and subcategorial information emerge simultaneously and conditioned by the available categorial frameworks.

Though it is possible to trace the distinction between categorial and subcategorial information in the successive stages of cognitive development, as we shall see below, I should like to point out here that this distinction is not absolute, but relative to the cognitive development of the individual. Information that for one person is subcategorial may be categorial for someone with a finer-grained system of categorization. Ontogenetically, moreover, I would go so far as to say that unless we accept the existence of innate categories, then all sensory information starts by being subcategorial, and that it is only as the subject learns to group sensory information in categories or types of things, events and properties that a large part of that subcategorial sensory information becomes categorial. In fact, part of cognitive development consists in just that, in learning to categorize, to perceive categorial information contained in the sensory substrate provided by subcategorial information.

² Raffman, D. "On the persistence of phenomenology". In: Metzinger, T. (ed.) *Conscious Experience*, Schöningh / Imprint Academic, 1995.

Subcategorical information

According to Metzinger, subcategorical information is constituted by phenomenal contents governed by the principles of *presentationality*, *reality generation*, *non-intrinsicity* and *context-sensitivity* and *object formation*. Though I am not entirely in agreement with him as regards the validity of all these principles, they provide a convenient framework for the characterization of subcategorical information that I propose here.

a) The principle of presentationality

According to the principle of presentationality, subcategorical information is afforded by phenomenal content that is non categorial, of the present moment (linked to the presence), and not susceptible of being remembered and consequently impossible to identify at different points in time. It is sensory information of a specific kind, derived from the capacity of human beings (and many other organisms) to process sensory information concerning each of the senses, and potentially much richer and sensorially finer-grained than what is expressible in concepts.

When we concentrate our attention on the sensory aspect of things, we can appreciate shades of colour, nuances of form, aroma and taste, and a multitude of other sensory information for which we have no specific system of categorization. With regard to the perception of colour, for example, we have a system of categorization for basic colours (red, orange, yellow, green, blue, etc.), for some more particular shades (carmine, scarlet, salmon orange, ochre, lemon, turquoise, etc.) and for the “colourless” black and white, but our discriminatory capacity, though limited, allows us to distinguish many more shades for which we have no system of categorization, even though they form part of our phenomenal experience of colour. It is this kind of sensory information to which I try to refer with the expression “subcategorical sensory information”. As indicated above, it is information about here and now, tied to a subjectively experienced present, not susceptible of being remembered, and accordingly impossible to identify at different moments of time.

If we are asked what we had for breakfast this morning, we remember perfectly well that we had coffee and toast, for example, but we cannot remember the exact colour of either the coffee or the toast. Or suppose we enter a haberdasher’s to buy cotton thread of exactly the same colour as one of our favourite shirts, but without taking the shirt with us; clearly, if the shirt is red we will buy red cotton, not cotton of any other colour, but it is equally clear that when the haberdasher shows us the whole range of reds we shall be at a loss to decide which to choose, and shall quite likely be in for a nasty surprise

when we get home with the cotton we end up buying. If, on the other hand, we take the shirt with us to compare with the range of red cottons in the shop, we are perfectly capable of choosing the red that goes best with the shirt. In short, sensory information tied to the here and now is much more fine-grained and precise than what we can remember and, consequently, identify at different moments in time. And what goes for sight goes for the other senses, too. Is the sound I have just heard exactly the same as I heard yesterday, or the taste of the wine I have just drunk the same as that of the wine I drank a few days ago in this or that restaurant?

These examples, and a multitude of others that we can all imagine, highlight an important feature of our sensory information: that our perceptual capacity for discrimination is much more fine-grained than our perceptual memory. This is why I say that subcategorical sensory information that is perceived but not remembered is non-cognitive. If someone claims to know what a robin is but then mistakes sparrows for robins, we would conclude that he or she does not really know what a robin is. For a certain kind of information to count as cognitive, it must not only be perceived but also remembered in a way that allows identification at different moments in time.

In the course of development it is evidently possible to establish increasingly finer-grained systems of categorization that take into account increasingly subtle differences; and this is indeed what happens, to the extent that it is called for by the practical necessities of our interaction with our environment. A shepherd distinguishes among different kinds of sheep, a winegrower among different kinds of vine, a skier among different kinds of snow; but the layman sees only sheep, vines and snow where these specialists see different kinds of sheep, vine and snow. If at time t_1 a specialist is shown a certain kind or category of sheep, and at time t_2 a different kind, he can see without doubt that they *are* different kinds of sheep, however similar they may be in size, colour and other physical features that are not relevant to their categorization. The non-specialist, who does not pay attention to the sensory features that differentiate the various kinds or categories of sheep, may have serious doubts about whether the sheep shown at t_1 and t_2 are of the same or different kinds. I am inclined to believe - and there are neurophysiological grounds for this belief - that the process of categorization plays a fundamental role in the storage and retrieval of information in our brain.

To the extent that we have a system of categorization for the basic colours, objects that differ in having different basic colours are reliably remembered and identified as different. If their difference in colour has not been incorporated in a system of categorization, it may be a difference that is cognitively irrelevant, and consequently will not be remembered. And what holds for colours holds equally for other sensory information. We remember sensory information to the extent that it is relevant to our system of categorization, even though our

capacity for perceptual sensory discrimination is potentially much greater and can provide much more sensory information than is generally employed in the processes of categorization.

b) The principle of reality generation

The second feature of phenomenal content that Metzinger discusses as a principle, and one that is closely related to its presentational nature, is its existential nature. According to the principle of reality generation, phenomenal content or subcategorical sensory information (which is essentially the same thing) invariably functions in all standard situations as an existential quantifier for cognitive systems like ours. The mere sensory presence of subcategorical phenomenal content forces our cognitive system to assume the real existence of what is presenting itself to us in this way, regardless of any perceptual identification that we may make of it. In Metzinger's words, "it is precisely this architectural feature of the human system of conscious information processing which leads to the phenomenal presence of a world" and underlies its ontological constitution (Metzinger, 2003, p. 99).

Here I disagree. I believe that the existential factor is not subcategorical phenomenal but categorial content. Suppose someone experiences the sensory impression - subcategorical information - of seeing something like a mouse flit along the wall at the end of the long room in which he happens to be sitting. He gets up, goes around the room looking carefully, but sees neither a mouse nor anything like one; and in the course of this inspection he finds that there is no hole or chink through which a mouse could have escaped from the room. His conclusion? Obviously, that his impression was false, some kind of optical illusion. In other words, that his subcategorical sensory information did not force him to assume existence. If in the course of his inspection he had found a mouse crouching in a corner of the room - if his earlier experience had been supported by an experience affording categorial content - then this latter experience *would* have forced him to assume the existence of the mouse.

Actually, not even all categorial perceptual experiences imply existence, only the veridical ones. What a veridical perceptual experience is has been analysed and characterized in Chapter III of my book *Mente y mundo. Aproximación neurológica*, to which the reader is referred; here I shall limit myself to stating that the perceptual identification of some item is veridical if and only if it is intersubjectively compatible with the other perceptual identifications of the same item (Vázquez, 2007, Ch. III). If I introduce half a straight stick in a tank of water, why do I qualify as false my experience of perceiving it as bent? Simply because if I introduce my hand in the water and feel the stick I perceive it as straight; because the stick is experienced as more consistent than water and the

“bend” is located just at the surface of the water; because the bend disappears as I take the stick out of the water; etc. Since all these perceptual experiences inform me that the stick is not bent, I discard as non-veridical the perceptual experience that informed me otherwise. Furthermore, as is also discussed in Chapter III of *Mente y mundo*, in cases of non-veridical perceptual experience, whether erroneous perceptual identification, illusion or hallucination, we generally have at our disposal an explanation for why we were deceived.

c) The principle of nonintrinsicity and context-sensitivity

According to this principle, subcategorical sensory information is not constituted by an invariant core of sensory experience of a specific kind, as has been postulated by some empiricists in their attempts to base knowledge upon sensory information; on the contrary, it varies, depending on the context in which the sensory stimulation takes place.

As we saw in Chapter II, the sensory information we become aware of in the course of perception is the result of a long neurophysiological process involving the stimulation effected, the architecture of the perceptual system, and prior knowledge. Factors such as those associated with prior knowledge (available categorial systems) may possibly be more important at the categorial level than the subcategorical, but both categorial and subcategorical information depend on both the stimulation effected and the processing to which it is subjected by the architecture of the perceptual system. The “same” stimuli can give rise to different sensory experiences depending on which sensory receptors they are received by, the detail with which they are processed, and the context in which all this takes place. In the case of sight, if the stimuli are received by the cones of the fovea, it can give rise to sensory experiences affording very precise information about colours, shapes, spatial locations, etc.; whereas if they are received by peripheral retinal photoreceptors, or their processing is deficient because they are afforded insufficient attention, the resulting sensory information will be imprecise.

Suppose I want to thread a needle. To do so, I must centre my gaze on its eye in order for the stimuli arriving from that part of the needle to be received by the cones of the retinal fovea, the region of the retina that, as we saw in Chapter II, congregates a large number of ganglion cells with tiny receptive fields. As a result, the stimuli will be processed at high resolution, allowing me to discern the edges of the eye of the needle perfectly sharply and so to thread it. If the same stimuli are received by the photoreceptors of the retinal periphery, as would happen if I shifted my gaze from the eye of the needle to its point, they will be processed at low resolution in a way that does not capture fine details of the visual field such as the edges of the eye of the needle. Consequently, if my

gaze is focused on the pointed end of the needle I will not be able to insert the thread in the eye at its other end. The farther from the centre of gaze a ganglion cell is, the larger is its receptive field, until finally, hundreds of photoreceptors converge on a single ganglion cell, and the processing of the stimulation they receive accordingly affords only relatively imprecise information.

Additionally, as the principle of non - intrinsicity and context - sensitivity announces, subcategorical sensory information is sensitive to context. In other words, the same stimulation affords different subcategorical sensory information depending on the context in which it occurs. This is highlighted by the following examples of perceptual illusion.

In the classical Müller-Lyer illusion (Figure 10), the two lines are the same length but are experienced in the process of perception as having different lengths, due to the context provided by the angles at their ends. In the same way, in the Ponzo illusion (Figure 11), the two horizontal white rectangles are the same length, but the context provided by the railway lines makes us see the “farther” rectangle as larger. A similar size effect is produced in the Ebbinghaus illusion (Figure 12); the two central discs are the same size, but the context provided by the outer discs, which are larger in one case than in the other, makes our sensory experience one of seeing central circles of different sizes.

Figures 13 and 14 likewise reveal the influence of context in the processing of stimuli, but in regard to a different aspect of




Figure 10 Müller - Lyer illusion



Figure 11 Ponzo illusion




Figure 12 Ebbinghaus illusion

sensory experience, colour vision. In Figure 13 the two central rectangles are the same shade of green, but due to their different contexts our perceptual experience is one of seeing two different shades. Similarly, in Figure 14, squares A and B are the same shade of grey, but the context makes us experience them as being of different shades.

The influence of context in the processing of stimulation is such that it can't only force the same stimulation to be processed in ways that give rise to different sensory experiences, as in the above cases, but can also "fill in" gaps in stimulation (e.g. the blind spot caused by the optic disc), and give rise to a sensory experience of perceiving something that does not exist in the real world, such as Kanizsa's triangle, the "perceived" boundaries and greater whiteness of which are just a contextual side-effect of our processing the stimuli provided by the three notched black discs at its vertices (Figure 15).

To sum up, since both categorial and subcategorical sensory information result from the processing to which the architecture of the perceptual system subjects received stimuli, errors in this processing can mislead us not only at the level of categorial perceptual identification, but also at the level of subcategorical sensory information. It is accordingly a great mistake to try to establish subcategorical sensory information ("qualia", "sense data", "sensations", "sensible ideas", or whatever) as a kind of absolute empirical basis for our knowledge.



Figure 13 Contrast color effects




Figure 14 Contrast color effects




Figure 15 Kanizsa's triangle

d) The principle of object formation

Briefly, the principle of object formation is that conscious phenomenal content emerges simultaneously and interdependently with the processes of object formation. That is, in the conscious perceptual processes of adults already endowed with a categorical system subcategorical information emerges simultaneously and interdependently with categorical information. Because we identify what we are perceiving as a member of a certain category, we pay attention to the impressions or sensory features that are pertinent to this identification; and because we pay attention to these impressions or sensory features, we make that identification and no other - all this, without forgetting that both the categorical identification of what is perceived and the pertinent impressions or sensory features are the result of a long process involving stimulation from the outside world, the architecture of the perceptual system, and available categorical frameworks.

The simple example of Rubin's vase (Figure 3) illustrates how categorical and subcategorical information are mutually complementary in the process of categorical perceptual identification. If we concentrate our gaze on the central part of this figure, at the neurophysiological level we are ensuring that the stimuli from this part are received by the photoreceptors of the retinal fovea (the cones), while stimuli from the rest of the visual field are received by the receptors of the peripheral retina (mainly rods). Consequently, given the architecture of our perceptual system, the stimuli from the central part of the figure are processed in much greater detail than those from the periphery, giving us high-resolution sensory information about the colour and shape of the central part and only low-resolution sensory information about the periphery. So much in regard to bottom-up processing of information in the pathway from the retinal ganglion cells to the lateral geniculate nucleus to area V1 to the associative visual areas V2, V3, V4, V5, etc. But the brain of someone who is already endowed with a categorical system including the category "vase" analyses this sensory information in the light of its existing categorical system, with the result that, at the same time as he or she sees the colour and shape of the centre of the figure, this person also sees them as representing a vase; and because the figure is thus identified as a picture of a vase, attention is paid to the sensory information pertinent to this identification, and to none other. The categorical identification of the figure as representing a vase is influenced both by the subcategorical sensory information derived from the processing of stimuli, and by the system of categorization at the disposal of the perceiving subject. In the perceptual processes of adults endowed with a categorical system, information flows not only upwards but also downwards. The associative visual areas projects a multitude of axons towards the primary visual area (V1), and from this latter a multitude are projected to the lateral geniculate nucleus.

If a perceiving subject had no concept of vase allowing the category of vases

to be distinguished, he or she could still see the colour and shape of the centre of the figure - always supposing, of course, that for some reason they were paid due attention. But a vase would not be seen in this colour and shape. That we almost mechanically pay attention to the colour and form is precisely due to our seeing a vase in this sensory information.

Suppose now that, as is usually done, we are asked to pay attention to the white shape framing the vase. When we concentrate our gaze on this other aspect of the figure, it is the stimuli therefrom that are received by the cones of the retinal fovea and processed in greatest detail, while the shape and black colour of the vase are now deprioritized and processed in lower resolution. As a result, in the white shape framing the vase we now see two heads facing each other.

In this example, in which the figure presented is ambiguous, it may seem that the subcategorical information (shape and colour) is perceived before the categorial information, at least in the case of the less easily perceived categorial content, the two heads. But this does not happen in unambiguous cases, in which, in accordance with the principle of object formation, conscious phenomenal content emerges simultaneously and interdependently with the categorial identification of what is being perceived.

Because I identify something as a bird, I pay attention to the sensory features that allow me to identify it as such; and because I pay attention to these sensory features, that is the categorial identification I make, none other. Someone whose categorial system includes not only the category of bird but also the subcategories of sparrow and robin, upon seeing a sparrow or a robin pays attention not only to the sensory features that allow its identification as a bird, but also to those that define it as a sparrow or a robin; and because he or she pays attention to these features, they are processed in such a way that the subcategorical sensory information they provide is precise enough to allow a member of the appropriate category (a sparrow or a robin) to be seen in them. Those who have no concepts of sparrow and robin that allow them to distinguish these two categories, upon seeing one of these birds, will most likely pay attention only to features that allow it to be seen as a bird (something with two wings, two legs, and a beak). Features that allow its identification as sparrow or robin, but are irrelevant for its identification as a bird, will be paid no attention; and because they are paid no attention, because they are irrelevant for categorial identification as a bird, they will likely be processed so deficiently that they go practically unperceived.

Categorial information

The foregoing brief discussion of the principle of object formation has to some extent anticipated what I wish to say here concerning categorial information. As was asserted at the beginning of this chapter, categorial information is constituted

by content that is cognitively as well as attentionally and behaviourally available; content that is not inexorably tied to the here and now, but can also be identified at different moments in time. In regard to categories with a specific sensory modality, categorial information includes, for example, information on basic colours, shapes, and well-defined tastes and aromas. The colours red, yellow, green, blue or brown are not only identified as such at the time of perceiving them; we retain a memory of that colour category, and can identify those colours as such at any other time. If we are shown a red card and asked to pay attention to its colour, after which it is shuffled with six or seven other cards of different colours, we shall evidently have no problem in recognizing it among the others. The same goes for categorized shapes (triangle, square, rectangle) or sounds (birdsong, the bark of a dog, the howl of a wolf). As members of a certain category of shapes or sounds, we have no difficulty in distinguishing them from the members of other categories. And although all these cases, and others that could be mentioned in relation to each of the senses, concern categorized sensory information of a specific sensory modality, categorial information can arise not only from the processing of unimodal sensory information but also from the processing of multimodal sensory information provided by more than one of the senses. Thus by “categorial information” I mean the sensory information that is obtained from stimulus-processing-derived sensory information by grouping in classes or types of thing, event, property, and so on, and that we cognitive beings store in memory in the form of categorial content.

Of the two types of sensory information, categorial and subcategorial, it is categorial information that allows us to make inferences, have expectations, identify types of thing, and hence have knowledge about our environment. It is therefore primarily categorial information that is important in the processes of perception; subcategorial information has little importance other than as the substrate from which categorial information is extracted. Indeed, categorial information does not specify the subcategorial information from which it arises: I can identify an animal as a dog because I hear it bark, because I see it half-hidden behind a hedge, or because it is approaching me with who knows what intentions. In all these cases I identify the animal as a dog, but in each case I base this identification on different subcategorial sensory information.

Since what is of primary importance is categorial information, both cognitively and as regards our immediate interaction with our environment, the subcategorial sensory information on which it is based may in most cases be subjected to quite low-powered processing. To pay more attention to subcategorial information than is required for the corresponding categorial identification amounts to useless and counterproductive effort. Only on certain quite special occasions, as in the course of a philosophical discussion on the greater richness of presentational phenomenal information, do we carry out that more detailed processing of the stimulation. On these occasions we do indeed obtain sensory information that is more fine-grained

than is expressible in concepts; but as we saw before, this subcategorical sensory information is tied to the present and impossible to remember or, consequently, to identify at different moments in time. On most occasions, the subcategorical information on which the categorial identification of an item is based can be fairly imprecise.

It is through categorial information that we achieve intersubjectivity. Subjects who share the same categorial system live in the same world, enjoying perceptual access to the “same” categorial identifications, even though the subcategorical basis for these categorial identifications is different for each individual subject (and also for any given subject at different times). I see the dog from one angle, you from another; the stimuli you receive are different from mine, and your subcategorical sensory information is accordingly likewise different from mine. But since we both possess the concept of a dog, both you and I identify the animal we see as a member of this category; and in this sense we have access to the same world, or at least to a world in which our shared categories are constituted by enough common elements as to make intersubjective communication possible.

Not only can the same categorial information be borne by different constellations of sensory information; it is also possible for the “same” stimulation to induce different perceptual identifications by different subjects or by a single subject at different times, depending on the aspect of reality to which attention is paid in each case, and on the available system of categorization. Both you and I can identify the same animal as a dog at one time, as a mastiff at another, or simply as an animal on a third occasion, so long as we possess the corresponding concepts; and similarly, the animal identified by one of us as a mastiff may be identified by another simply as a dog, because he does not possess the concept of mastiff. Our categorial identifications of a given item depend not only on stimuli from the outside world, but also on our system of categorization. Someone who has never learnt to distinguish sparrows from robins, but only possesses the concept of bird, can identify sparrows and robins as birds but cannot identify sparrows as sparrows or robins as robins, even though he or she receives stimuli appropriate to and sufficient for the performance of these identifications.

In conclusion, in this chapter we have seen the differences and interactions between categorial and subcategorical sensory information, and how categorial information emerges in the light of previously acquired knowledge (the available system of categorization) in the processes of perception. In the next chapter we shall delve further into the nature of categorial information.



V. Categorical perception



As was noted at the end of the previous chapter, subcategorial information can be much more fine-grained than categorial information expressible in concepts, but as regards the praxis of our interaction with our environment it is much less informative and interesting. Its presentational nature, tied to the here and now, means that only in very special situations is it worth paying attention to for its own sake. This is not to say that subcategorial sensory information is not present in every act of perception, but it is present only to the extent that it gives access to categorial information.

Let us imagine someone who, presented with a dog, is capable of extracting from the scene all the subcategorial sensory information that would be extracted by any normal person, but is incapable of further extracting any kind of categorial information. That is, this person detects, like any other, the colour, shape, spatial location and movements of the animal, and also the subcategorial sensory information provided by the other senses (its smell, the sounds it makes when it barks, moves and breathes, etc.), but without all this sensory information inducing any categorial information: he or she fails to identify, in the observed shape, the shape of a dog, or even the shape of an animal; fails to identify, in its movements, the movements of a live being; fails to identify, in its bark, the barking of a dog rather than just a noise; and so on. Even though he or she extracts all the subcategorial sensory information that can be extracted by processing the sensory stimuli from the animal, this person is as devoid of information about the dog as a child is before learning to categorize, to identify types of thing. For this person, all the sensory information is new, and is again new every time the dog crosses his path. And if the same is true not only of information concerning the dog, but also of information concerning all the other objects and properties in the environment, that person would be, at all times and in respect of all objects in the environment, in a similar situation to that of a precategorial child. If there is a difference between a newborn child and a precategorial child aged 5 or 6 months, it lies in the almost complete inability of the neonate to see the world. Before the neonate can access subcategorial sensory information about the world, its neurons must learn to transform the stimulation from the environment to sensory information, something that in the case of colour perception does not usually occur before age 5 or 6 months.

The fictitious case of agnosia that I have just described is indeed an extreme case, because all the sensory modalities are affected; but phenomena that from a cognitive point of view are similar do actually occur in cases of perceptive visual agnosia, in which patients are unable to recognize objects by their shape

even though they may have normal visual acuity. They cannot identify what they see, do not know what objects they are seeing or of what use they are. Persons with perceptive visual agnosia can draw the objects they see almost as accurately as a normal person; what they cannot do is to recognize those forms, the form of a house, or a bird, or a tree, say. In *To see but not to see: a case study of visual agnosia*, Humphreys and Riddoch describe the case of a patient who suffered a stroke following a surgical operation; a blood clot reached his brain, blocking the arteries that feed the occipital lobe. Afterwards, he was unable to recognize objects in his surroundings even though he clearly saw their shapes, as was evidenced by his being able to copy drawings quite accurately but not to recognize what he was drawing. When shown a drawing of a carrot with its leaves (Figure 16 a) he remarked “I haven’t the least idea what it is. The bottom end looks solid and the other has feathers. It doesn’t seem logical, unless it’s some kind of tree”; and upon being shown a drawing of an onion (Figure 16 b) he said “At this moment I’m absolutely lost It has sharp points on the bottom, like a fork. It might be a collar of some kind.” In terms of the distinction between categorial and subcategorial sensory information, this patient and many other agnostic persons fail to see the appropriate categorial information in their subcategorial sensory information. Their processing of visual sensory stimuli is not mediated by a categorial system.



Figure 16 Perceptive visual agnosia

A less serious deficiency is prosopagnosia, difficulty in recognizing particular faces. Confronted with a face, patients suffering from prosopagnosia recognize that it *is* a face, but are unable to recognize whose face it is, even when it belongs to a friend or a close relative such as a parent or sibling, in spite of being able to recognize these people by other means (e.g. by touch, from their voices, or from the sound of their footsteps).

Unlike patients suffering from perceptive visual agnosia, people with normal vision see types of things, events and properties in their sensory information. Our world is categorized. We see chairs in the sensory information about chairs, tables in the sensory information about tables, books in the sensory information about books, and so on for all the various types or categories of object in the world. When we go outdoors we see buildings, balconies and windows, men, women and children moving from place to place, and on seeing them we know effortlessly what they all are. Nor is this capacity for recognition specific to a single sensory modality; on the contrary, it holds for each of our senses. In sensory information derived from a barking dog I hear a dog barking, not a

mere noise lacking categorial identification; and the same goes for the sound of a doorbell, the song of a bird, a human voice, the taste of a wine, or the aroma of a perfume.

The perception of adult individuals depends not only on the stimuli received and the architecture of the perceptual system, but also on the categorial system at their disposal. In the process of perception, categorial information emerges interdependently and simultaneously with the sensory information obtained by processing stimuli arriving from the outside world or from one's own organism. From a cognitive point of view, as noted above, it is the categorial information that is important. Because I see a carrot in the form and colour of the root and leaves of a carrot, and an onion in those of an onion, my perception is categorial and informative. Because I identify that colour and form as the colour and form of a carrot or an onion, I can anticipate their internal structures, properties, culinary uses, etc. This knowledge is inherent to categorial information, not to the subcategorial sensory information on which it is based. Subcategorial information varies from one moment to another, it is not the same when I see a carrot or an onion half a metre away as when I see them 3 or 4 metres away, and it also varies with the position they occupy, with whether I have my glasses on, and so on; but in all these cases my categorial identifications of the carrot as a carrot and the onion as an onion are the same, and these identifications allow me to infer that they possess the sets of properties associated with these two categories of thing.

The subcategorial sensory information on which the corresponding categorial information is based is in some ways cognitively irrelevant. The subcategorial sensory information on which you base your categorial identification of the carrot or onion is different from the subcategorial sensory information on which I base my categorial identification of the same carrot or onion; but to the extent that we possess the "same" categorial system, we both identify the carrot as a carrot and the onion as an onion. And because we share the "same" categorial system and can make the "same" categorial identifications, we can communicate with each other. If I ask you to pass me an onion to add to the salad I am preparing, I know that what you are going to pass to me will be an onion, not a carrot. The objectivity and intersubjective validity of our knowledge stems from our sharing the "same" categorial system; because we share the same categorial system, we have perceptual access to the "same" world.

At the level of subcategorial sensory information, we live in the private worlds postulated by sense data theorists. But to the extent that subcategorial sensory information always bears categorial information, the world of objects and properties - the categorized world - is public. In fact, given the architectural similarity of one person's perceptual apparatus to another's, it may be supposed that even at the level of subcategorial sensory information there are no great

differences between the sensory experiences of different individuals. In general, we all have access to the same or very similar shades of colour, and detect the same or very similar differences of sound, smell, taste, etc.

If I have been placing the word “same” in quotation marks when applying it to categorial information, it has merely been due to the fact that categories and their mental representations, concepts, are subject to evolution and development. Just as my concept of water is not now exactly the same as that I had when I was 5 or 6 years old, neither is it exactly the same as a chemist’s. Concepts are subject to phylogenetic and ontogenetic evolution, and *a fortiori* so are categories, their referents. As individually and collectively we acquire more and more information about the things there are in the world, our concepts of these things become richer and more precise. Consequently, to say that two or more individuals have the same categorial system can mean no more than that many of the contents ascribed by these individuals to any given type or category of thing, event or property must be shared. Or perhaps, since it is not usually easy to specify shared content, it may be less problematic to say that two or more individuals share the same categorial system to the extent that they make the same perceptual identifications.

Just as different individuals, or the same individual at different times, can access the same categorial information on the basis of different stimuli, so too can it occur that the same stimuli can afford different categorial information to individuals whose categorial frameworks differ in relevant ways. A shepherd identifies different types or categories of sheep in a flock in which the layman only sees individually distinct sheep of a single type or category. Similarly, an expert in viticulture distinguishes at any time of year among different types or categories of grapevine: where the layman sees simply vines in a vineyard, the viticulturist sees albariño, garnacha, palomino, or whatever, i.e. the particular variety of grapevine that has been planted in that vineyard.

That different individuals, or the same individual in different situations, may make different categorial perceptual identifications does not mean that the categorial information with which the things of the world present themselves in the processes of perception is arbitrary or depends entirely on the perceiver’s categorial framework. If there were no sensory features allowing the shepherd to distinguish among different kinds of sheep and the viticulturist among different kinds or categories of grapevine, the corresponding categorial systems could never have been established and the experts would not have appropriate sensory information in which to see the categories they do see. It is because of the existence of those differences in sensory features that the corresponding categories can be established. Nevertheless, in some community it might well happen that, even given the existence of these sensory differences, they might be ignored as irrelevant to the interactions of the members of the community

with their environment, with the result that the corresponding categorial system would not be constituted. And in the absence of that categorial system, those sensory features would probably be paid no attention, and would consequently not even be perceived by the members of the community. To perceive an item with a certain degree of sharpness we must pay it attention, concentrating our gaze upon it, so that the stimuli therefrom are received by the cones of the retinal fovea.

Our categorial system leads us to pay attention precisely to the kind of sensory information in which a certain categorial information can be seen. Subcategorial sensory information that does not contribute to categorial identification is ignored, and even the subcategorial sensory information that does contribute to categorial identification is only made use of to the extent that it allows that identification. As I have already repeated several times, the information that is important in processes of perception is categorial; once categorial information has emerged, it matters little how fine-grained the subcategorial sensory information was upon which it was based.

Because I am thirsty I go to the fridge for a beer. The sensory information on which I base the identification of one of the cans of beer in the fridge may be quite diverse. The can may be in full view at the front of the fridge; it may be only partly visible; it may be completely hidden, so that I have to fumble for it among other drinks, identifying it by touch. But it does not matter what sensory information I base the identification of the can of beer on, that is irrelevant; what is relevant is that in each of these sets of sensory information there lies the categorial information of a can of beer. To pay undue attention to sensory information, over and above what is necessary to extract from it the corresponding categorial information, constitutes pointless effort. The sensory information on which I base my perception of a bear in the woods matters little; what is important is that I identify the animal as a bear, that this identification be veridical, and that I consequently adopt appropriate cautionary measures.

The more sensory information I have, and the more varied its sensory modalities, the more reliable will the corresponding categorial identification be. My identification of an orange from a distance of 2 or 3 metres is not as reliable as my identification of an orange I can touch, smell, peel and taste, as well as see. The more perceptual identifications we make of a given item, and the more varied, the more certain its identification will be.

In everyday practice, the sensory information on which our categorial identifications are based can vary greatly both qualitatively and quantitatively. Even so, the architecture of our perceptual system is generally remarkably efficient in the task of providing us with relatively reliable categorial information. At times it makes mistakes, but when it does it is usually either because we have been too hasty in identifying something, or because we have laid a kind of trap

for our perceptual system, as in the cases of the Müller-Lyer arrows, the rectangles on Ponzo's railway lines, and the vast majority of other perceptual illusions. Since both categorial and the subcategorial sensory information are the result of a long process involving stimuli, the architecture of the perceptual system and prior knowledge, it is always possible that a mistake may be made in the course of processing the stimuli, we can always make mistakes; but though perception can at times deceive us, it is also through perception that we recover from our errors. Visually, we identify the Müller-Lyer arrows or Ponzo's rectangles as being of different lengths, but it is through other perceptions of these arrows and rectangles that we find these apparent differences to have been the fruit of optical illusions.

I repeat: the categorial nature of the things of the world that we access through processes of perception is not something that exists independently of their *being* things; it is their *way* of being things for cognitive beings who possess the corresponding systems of categorization. And our access to things is direct. This direct access is evidently not to an assumed world of things in themselves, but to the world of things, events and properties that emerge from the processing, by the architecture of our perceptual system, of the stimulation received from those things, events and properties. What the world may be, other than in its manifestation of itself through the senses (alone or aided by observational instruments), is an issue that is senseless, because it has no possible answer.

In the processes of perception, the categorial nature of the things of the world presents itself to us as being out there before us, and our access to it, in the sense specified in the previous paragraph, is direct. To speak of the content of perception as representational is meaningless, an idea that, as we shall see in the next chapter, can only be entertained in the context of an untenable full-blooded metaphysical realism that postulates "things in themselves".


It is through the processes of perception that the most basic categories and the first concepts are constituted. Categories do not exist independently of things themselves, they are the way in which things, events and properties are things, events and properties in the world for cognitive beings. And our access to them in the processes of perception is direct. In the processes of perception, things, events, properties and so on *present* themselves to us, and it is concepts, i.e. the more or less hazy memory of these presentations, that are *representational*.

I see a friend in the street. My perceptual access to him is direct, in the act of seeing him there is no kind of intermediary interposed between us; my perception is my perception of my friend. Once he has passed, I have more or less hazy memories of this and other presentations of him. It is the



content of these presentations stored in my memory that is representational. This content, whatever it is, stands for something other than itself – in this case my friend. But in the instant in which I am perceptually identifying him, the content of my perceptual experience is my friend, as processed by the architecture of my perceptual apparatus.





VI. The presentational nature of perceptual content



Following on from the final paragraphs of the previous chapter, in the present chapter I show that the informative content of perception is not representational but presentational. It is concepts that are representational, that is, the more or less vague memories of perceptual content that are stored in memory. In short, the informative content of perception does not exist independently of the things, events, properties and states of things of the world; it is rather the way in which these things, events, properties and states of things manifest themselves to cognitive beings in the processes of perception.

In the precise moment in which I am seeing and touching the keyboard of my computer, it is only one thing that presents itself to my consciousness, the keyboard, not a representation of the keyboard, or any such thing. What happens when I leave my study and the keyboard is no longer present is a different kettle of fish. If I think of the keyboard after leaving my study, what is activated in my memory is the more or less vague recollection of that previously perceived categorial content. It is this recollection that is representational, since it stands for something other than itself, *viz.*, the keyboard. But, I repeat, when I am perceiving the keyboard, the only thing that presents itself to my consciousness is the keyboard, not any kind of intermediate representation of it. It is in this sense that we can also say that our access to the perceived world is direct. What we must not lose sight of is that this direct access is not access to an hypothetical world of “things in themselves”, but to the world as it presents itself to us in the processes of perception upon being processed by the architecture of our perceptual system.

To speak of representation assumes the existence of two different things, the representation and what it represents. It is the existence of these two things that constitutes one of the fundamental differences between the content of beliefs and memories on the one hand, and the content of perception on the other. In the case of beliefs and memories, the content of the belief or memory is one thing, and what it is a belief or memory of is something else. I may believe that tomorrow it will rain, because that is what the weatherman has just said, but my belief is one thing and whether it rains tomorrow is another. In the same way, I may remember that yesterday afternoon, on leaving work, I forgot to turn the computer off; but in this case, too, the content of my memory is one thing and whether I left the computer on is something else.

With perception it is quite different. When I am seeing and touching the keyboard of the computer, only one thing presents itself to my consciousness, the keyboard. If I set about describing the content of my perceptual experience,

I would be describing the keyboard, just as it presents itself to me in the act of perception; I would not be describing any representation of the keyboard.

A metaphysical realist might argue that our not being conscious, at the time of perception, of both the representation and what is represented, does not mean that they do not exist, since our perceptual experience of the things of the world is one thing and things of the world as they are “in themselves” is another. But this reply, the fruit of untenable dualist prejudice, just does not hold water. If the content of our perceptual experience were representational, what would it represent? A world we have never and will never have access to?

Even supposing, like the metaphysical realist, the existence of an independent outside world, endowed with certain properties that it is assumed to possess but which are not the properties that we cognitive beings can access, it is inadequate, to say the least, to characterize the content of perception as representational of we do not know what. A representation *is* a representation to the extent that it stands for something else, and we call it representational because we have access both to the representation and to what it represents, what it stands for. Only in so far as we have access to both these things, the representation and what it represents, are we authorized to describe one of them as a representation of the other.

A map of a city or a region are representations of that city or that region because there is a correspondence between the map and the city, or between the map and the region. We use the map because tracing a trajectory on it allows us to predict the results of our movements in the city or region. Similarly, even though the chemical formula of a molecule and the mathematical formula for the trajectory of a projectile bear no physical resemblance to the molecule or the trajectory of the projectile, these formulae are representational because they allow us to predict the behaviour of the molecule and the projectile. But in the case of perception, what is it about the hypothetical independent world-in-itself that we predict by examining the content of our perceptions? What is the correspondence between the supposedly representational content of our perceptions and the world-in-itself, and what use can we make of any such correspondence? How can I decide whether the supposed representation is correct or not if I only have access to the representation and not to what is represented?

To think of the informative content of perception as representational is one of the unhappy consequences that inevitably await the full-blooded metaphysical realist. If we incur in the fallacy of placing mind and world in two different closed compartments, as the metaphysical realist does, we are left with no alternative for the content of our perceptual experience but to join with representationalists in considering perception as representation. One possible way of avoiding perceptual representationalism is argued for by the small

group of direct realists; but in the context of full-blooded metaphysical realism this avenue has the serious drawback of being unable to provide a satisfactory explanation of illusory and hallucinatory perceptions, as we shall now see.

It is clear that in the cases of both illusory perceptions and hallucinations we cannot speak of direct access to the world of “things in themselves”, since the content of the illusory or hallucinatory perceptions is one thing and reality another. In the case of illusory perceptions the perceived properties are not the properties really possessed by the objects of perception. For example, the difference in length between the two arrows that is experienced in the process of perceiving the Müller-Lyer figure is not a property of this figure in the real world. And in the case of hallucinations, “the perceived” object may not even exist in the real world.

A second problem that faces direct realism, within the framework of metaphysical realism, is derived from the perception of secondary qualities (colour, aroma, heat, cold, etc.), since these properties, as experienced in the processes of perception, are not exactly those supposed to be possessed by real-world objects.

Issues like those I have just mentioned have led the representationalist majority to distinguish between the content of our perceptual experience on the one hand, and the independent world of “things in themselves” on the other. But in so doing they place themselves in the position of having to face problems that are different but no less thorny.

The representationalist alternative easily surmounts the difficulties posed by illusory perceptions, hallucinations and secondary qualities for the direct realist, but only at a cost: the introduction of an intermediary, the representation, which acts as a “veil of perception” between mind and world. And the cost is heavy: representationalists find themselves in the awkward position of being unable to make sense of their notion of representation, since all we mortals have access to, according to representationalists, are the representations themselves, not what they are supposed to be representations of. To speak of a representation of something we cannot know, because it cannot be accessed, is meaningless.

If, on the contrary, we adopt the epistemological framework provided by the analysis of perception pursued here, then like the naïve realist we can claim that our access to the world is direct, but with the difference that this world is not the metaphysical realist’s hypothetical world of “things in themselves”, but the world as it manifests itself to cognitive beings in the processes of perception, the world resulting from the processing to which the perceptual system subjects stimuli arriving at its sensory receptors. To ask what the world may be, other than as it manifests itself to cognitive beings in the processes of perception or experimental observation, is a meaningless question because it has no possible answer.

It is the processed world of our experience that can meaningfully be said to be real, to be “out there” before us, and to be independent of what we can say or think about it. We can say that it is independent of what we say or think because, unlike the content of mind, its existence does not depend on our will. He who perceives receives something from without, whereas concepts can be endogenous and arise spontaneously. While I write at my computer, I identify its keyboard visually and by touch. If I stop looking at it, my visual perception of it disappears but I am still conscious of its existence through touch. When I leave my workplace, someone else can go on using it; and every morning, when I come back to my workplace, it is still there. These facts, and many others that could be adduced, are the proof of its independent existence. And we may say that this independent existence, with all its properties of colour, shape, size, etc., is something possessed by the keyboard (and similarly for objects in general); but to say this can only mean that the properties of the keyboard (or of the object in question) can be experienced by a cognitive being, whether directly or indirectly with the aid of instruments.

Consider colour. According to physics, colours do not exist, they are not a property of material objects. When an object is illuminated with white light, it reflects light waves; it is the waves that really exist. It was arguments of this kind that led Locke and certain other philosophers to distinguish between primary and secondary qualities. My view is quite different. Granted, when an object is illuminated with white light, it reflects light waves; in the dark, there are indeed no colours. But that light, or as I would prefer to put it, the stimulation from the objects illuminated with that white light, when received by the cones of the retina and processed by the architecture of our perceptual system, is experienced by human cognitive beings as a sensation of colour or, more exactly, as the property of colour with which the illuminated objects present themselves to us in the process of perception.

The property of colour belongs to objects as perceived by human beings, and is accordingly one of their ways of being objects for human cognitive beings. If the architecture of our perceptual system were not endowed with cone cells or some equivalent organ, colour would not be a property of things. But if we were not endowed with cone cells or some functionally equivalent type of sensory receptor, neither would shape, a primary quality, be a property of things, because we should not perceive the differences in shape that allow us to distinguish between some things and others; and the same could be said of many other properties. And if we had no spectrophotometers or equivalent instruments, wavelengths would not form part of the entities postulated by physics. In short, colours are real, with a reality status equivalent to what can be assigned to wavelengths.

If I want to decorate a room, the information about the real world that colours

provide me with is better than that provided by wavelengths. But if my purpose is to distinguish finer differences in colour than are inaccessible to the senses, then the information provided by the spectrophotometer is evidently preferable.

To ask what there is, or to say what is real and what is not, only makes sense in regard to the things of the world in relation to cognitive beings and the instruments they use in observational processes. Certainly, the noise made by a tree falling in a wood exists whether or not there is a cognitive being there to hear it; but the only thing it can make sense for this to mean is that if a cognitive being had been there, it could have processed the stimulation emitted by the falling tree and identified the result of that processing as a noise. And what goes for the noise goes equally for many other effects of the falling tree.

Thus it is the content of the world as processed by the architecture of our perceptual system that we describe as presentational rather than representational. In acts of perception the world presents itself to us, we have direct access to it, as claimed by common-sense realists. But, as I have remarked so often, this direct access is not to a world of “things in themselves”, but to the world processed by our perceptual system; and it is the memory of these presentations that is representational.

When perception is taking place, its informative content is presentational; and its being presentational is nothing other than the way the things of the world *are* things of the world for cognitive beings. Once anything (any thing-of-the-world) is no longer present (no longer presents itself to us) it is the memory of its presentations that is representational, and which, being representational, stands for something different from itself. But what our most basic mental representations of the world stand for is not the metaphysical realist’s hypothetical world of “things in themselves”, but the world of our perceptual experience. Thus understood, the word “representation” makes sense, but the representers are not the informative content of perception, but its memories, mental representations. It is concepts, the mental representations of categorial informative content, that we may meaningfully describe as representational, regardless of the uncertainty about their nature that was discussed in Chapter III.

The analysis of perception I have been pursuing reveals a structure that is quite different from that portrayed by extreme metaphysical realism. Instead of a realism that, as Putnam would say, commits the fallacy of division, placing the mind and the world in two watertight compartments without any possibility of communication between them, the new structure not only includes that communication, but is also free from the problems encountered by direct realism and representationalism.

Since our access to the world through perception is access to the world afforded by the processing to which the architecture of our perceptual system subjects the stimuli received from entities perceived, there is always a possibility

that errors of processing occur, but this does not pose insuperable problems for the account of perception that is being proposed here. Consider a perceptual illusion such as that affecting visual perception of the Müller-Lyer arrows. The visually perceived relative lengths of the two shafts are certainly not those they actually have in the real world. In what real world? The world of “things in themselves”? Not at all. We say that the perceived relative lengths are not those the arrow shafts actually have in the real world, and that we are therefore experiencing a perceptual illusion, because the perceived relative lengths are not compatible with those manifested by these arrow shafts in all the other perceptual experiences we can have of them: for example, when we measure the lengths of the shafts, when we superimpose one on the other, when we eliminate the context provided by the arrowheads and flights, etc. All these other perceptual experiences of the Müller-Lyer arrows provide us with information on their relative lengths that is incompatible with the relative lengths manifested in direct visual perception of the unmanipulated arrows; consequently, we regard as real the relative length assessed through these other perceptual experiences, which we deem more reliable. And we judge that our direct visual perception is not veridical, because in this case the visual context causes the architecture of our perceptual system to err in processing the stimulation received.

Our access to the world is direct, but we only attribute reality to the world to which access is gained through veridical perceptual experiences. And as we have just seen, we deem the perceptual experience of an item to be veridical if it is compatible with the rest of the perceptual experiences of that item, and we deem it to be non-veridical if it is incompatible with those experiences (Vázquez, 2007, Ch. III). In the end it is always of the world of our experience, the world of our veridical perceptual experiences, that we predicate reality. As I noted previously, perception sometimes deceives us, but it is also through perception that we recover from our error.

The above analysis of the visual perceptual illusion caused by the Müller-Lyer arrows is equally applicable to all other perceptual illusions, regardless of how they are caused or their sensory modality.

The account of perception that I am proposing allows a very similar analysis of hallucinations. The only difference is that in the case of hallucinations our perception has no external referent, even though the person suffering the hallucination experiences the content of his or her perception as if it were real. Just as perceptual illusions occur due to distorted processing of stimuli, hallucinations are generally due to faults in the architecture or working of the perceptual system, usually in the brain. But just as in the case of perceptual illusions, we deem the informative content of a hallucination not to be real because it is incompatible with the remainder of the subject’s perceptual experiences, or intersubjectively incompatible with the perceptual experiences of other subjects.

That is, we deem the content of a perception to be hallucinatory because it is subjectively or intersubjectively incompatible with the perceptual content of the remainder of our experiences. So in the case of hallucinations, too, it is the world of our veridical perceptual experiences that acts as the criterion of reality, not an hypothetical worlds of “things in themselves”. If we deny reality to the world of hallucinations, it is because that world is incompatible with the real world that we access through veridical perceptual experiences.

Given the possibility of error in the perceptual processing of stimuli, and given the distinction between veridical and non-veridical perceptual experiences, it is clear that perceptual illusions and hallucinations no longer pose the problems they posed for direct realists, any more than do the perception of colour and other secondary qualities, which we have already seen to pose no problems for my account of perception. I stress, moreover, that the above explanation of these two types of non-veridical perceptual experience – illusions and hallucinations – is not only compatible with the presentational nature of the informative content of veridical perceptions, but in fact presupposes that presentational nature, thus making it unnecessary to suppose the existence of an intermediary (the representation) that mediates between the mind and the world.





VII. Categorical perception and the empirical justification of beliefs



Unlike the many anthropologists and cognitive psychologists who are in the habit of using the terms “category” and “concept” as synonyms, I have here followed Murphy (2004) and Prinz (2004) in giving them different, albeit closely interrelated, uses. By “category” I have understood the result of organizing sensory information derived from stimulus processing in meaningful classes on the basis of common features; and I have stressed that, so understood, categories are not independent of things themselves, or of the way in which things manifest themselves to cognitive beings in the processes of perception. The content of a category is thus presentational, though susceptible of being stored in long-term memory and identified transtemporally - hence its categorial nature: to identify something perceptually as a certain kind of thing (as a tree, a flower, a bird, a colour, etc.) is to identify this categorial nature. The term “concept”, on the other hand, has been used here to refer to the mental representation of categories, which in the case of human beings is fundamentally symbolic. In terms of Frege’s distinction between sense and reference (though without subscribing to his ontology, the ontological setting being what we have seen in previous chapters), we might say that the concept expressed by a sign is the sense of the sign, and the category its reference.

I now propose to provide a criterion of truth for perceptual statements - and only for perceptual statements, which due to their proximity to experience are the only statements that can be compared directly with the categorial content of our perceptual experience. Such comparison is possible because the concepts involved in the constitution of perceptual statements are the mental representations (and in this case also the symbolic representations) of the categorial content of that experience: among the many different kinds of concept, with very different degrees of abstraction, perceptual statements are characterized by their descriptive terms all referring to perceptual categorial content. In perception, categorial content presents itself together with the sensory information on which it is based, whereas concepts are the mental representation of that categorial content, and perceptual statements fulfil the extraordinary function of giving symbolic expression to this conceptual content. Thus perceptual statements, and hence also beliefs expressed by this kind of statements, can find empirical justification in the corresponding perceptual experiences.

If someone emits a perceptual judgement, such as “robins have an orange-red breast and face”, the descriptive terms employed refer to categories that can be objects of perception. “Robin”, “breast”, “face” and “orange-red” all refer to categories that can be perceptually identified in the corresponding processes

of perception, and the perceptual judgement “robins have an orange-red breast and face” can therefore be verified empirically by the corresponding acts of perception. It is this that authorizes us to say that beliefs expressed by this kind of statement find their empirical justification in perception - so long as it is recognized that, in the context of the theory of perception presented in the foregoing chapters, to speak of justification does not imply holding that beliefs expressible by perceptual statements are incorrigible, and thus the ultimate foundation of empirical knowledge.

As we saw earlier, the present account of perception allows the possibility of distinguishing between veridical and non-veridical perceptual experiences. But this implies that not even experiences deemed veridical provide an ultimate foundation for empirical knowledge, if we demand of an ultimate foundation that it be incorrigible. There is always a possibility that perceptual experiences hitherto deemed veridical, and hence as providing a sure basis for our empirical knowledge, may not in fact be so. The following two examples are illustrative.

For generations, it was accepted that two far-away events either are or are not simultaneous, regardless of the state of motion of any observer; but since the emergence of the theory of relativity in the early 20th century we know that this is not so. In other words, certain perceptual identifications that were considered veridical until 1905, the year in which Einstein published the special theory of relativity, ceased to be considered as such.

Like the above, the following example can be used to show the inexistence of incorrigible perceptual experiences. Suppose that on leaving home one morning I come across the neighbour with whom, the previous day, I had been discussing a problem of the community association. I try to pick up the conversation where we left off, but instead of following me my neighbour looks puzzled and protests that he doesn't know what I'm talking about. I am taken aback, and try to remind him of yesterday's conversation, but after listening to me for a while he provides the answer to this little mystery: “You're mistaking me for my twin brother Andrew”. I apologize, “Sorry, I didn't know Andrew had a twin brother”; and in so doing I am acknowledging that my perceptual identification of this man as my neighbour Andrew was non-veridical.

As the above examples show, there is always a chance that we may be mistaken in considering as veridical the information provided by a given perceptual experience. However, it so happens that in the vast majority of cases our perceptual experience provides us with information that we regard as veridical and have no reason to question. It is because of this that I hold veridical perceptual experience to provide an empirical justification of perceptual statements, a justification that in everyday life we consider to be not only sufficient but one of the most solid justifications of our beliefs.

Perceptual judgements are similar to the observational statements of

logical empiricism and the basic statements of Popper, but without either the incorrigibility attributed to observational statements by logical empiricism or the conventional character of Popper's basic statements. They are akin to the former in being the empirical foundation on which our knowledge is based, and akin to the latter in that different communities of speakers may have partially different categorial systems. However, in the present account of perception conventionalism has a limit marked by the establishment of the categorial systems themselves: if we can organize the things of the world in categories, it is because there exist sensory differences, or differences of use, that allow environmental stimuli to be organized in meaningful classes - otherwise, it would be impossible for different categories to be established. This does not prevent different communities from developing different categorial systems, depending on whether their attention is focused on one or another type of sensory features.

With these caveats, and in keeping with the story so far, I propose the following truth criterion for perceptual statements (and only for these):

A perceptual statement p belonging to a language L is true in relation to the categorial world M (and is therefore empirically justified) if and only if there exist veridical possible perceptual experiences x_1, \dots, x_n such that p expresses the categorial content of x_1, \dots, x_n .

At first sight it may seem that this proposed criterion of truth or empirical justification is circular, because of the appearance of the term "veridical" on the right-hand side of the biconditional. It is not circular, however, because the criterion for characterizing the categorial content of a perceptual experience as veridical is established at the level of perceptual experiences themselves, at the level of presentational content, whereas the truth is predicated of the bearers of truth, in this case perceptual statements, at a different, representational level. The criterion is thus a correspondence criterion of truth; but unlike the classical correspondence criterion, it examines correspondence between two experiences of the subject, one representational and the other presentational. It therefore avoids the kind of "fallacy of division" that Putnam warned of in the Preface of *Realism with a Human Face*, a fallacy in which the classical correspondence theory of truth does incur in its assignment of truth-bearers and of the empirical reality by which they should be justified to levels that are not only different but mutually incommensurable.

I emphasize that the above criterion of truth concerns only perceptual statements, and that any application concerns only the perceptual statements of a given language L and the categorized world M of which L is the symbolic representation. The reasons for these restrictions are evident. Although many categories can be shared by multiple languages, and many are in fact so shared, so that the "same" propositional content can be expressed in different languages,

it is through their linguistic expression in a given language that propositional contents present themselves to the users of that language. The proposed truth criterion accordingly refers to perceptual statements of an arbitrary given language. Furthermore, however, certain categories are or may be specific to a particular language or to a particular specialization within a language, so that the truth of a perceptual statement p in language L can only meaningfully be spoken of in regard to the speakers of L who possess both the set of concepts used in the formulation of p and the corresponding perceptual categories. The perceptual statement “The bird in the cage is a goldfinch” makes sense solely for individuals who possess the concepts of “bird”, “cage” and “goldfinch”; those who do not possess the concept and category of “goldfinch”, but do possess the concepts and categories of “bird” and “cage”, can verify that there is a bird in the cage, but not that it is a goldfinch. To be able to verify that the bird in the cage is a goldfinch they must possess the concept and category of “goldfinch” as well as the concepts and categories of “bird” and “cage”. If they do not possess the concept and category of “goldfinch”, they cannot make full sense of the statement or determine whether the bird in the cage is a goldfinch or not. In short, to be able to determine the truth value of a perceptual statement p , it is necessary to be in possession of the categories expressed by the perceptual statement in question and the conceptual and symbolic representation of these categories; hence the relativization of the truth criterion to an arbitrary given language L and the categorial world M of which L is the linguistic codification.

To restrict the truth criterion not only to a language L and the corresponding world M but also to the categorial and linguistic frameworks of individuals does not imply solipsism. As was argued in Chapters IV and V, one of the features that differentiate categorial information from subcategorial information is its public nature. Subjects that share the “same” categorial framework live in the same world, and it is in relation to that shared world and to the truth-bearers encoding its symbolic representation (in this case the language) that the word “truth” has meaning.

According to the proposed truth criterion, not every perceptual experience justifies beliefs, only those that are veridical. The recognition of a perception as non-veridical automatically implies recognition that its content is not real, and hence that it is not legitimate to base the corresponding belief upon it. If I see that the window is open and think that my perception is veridical, then I am justified in believing that the window is open. By contrast, if I see the Müller-Lyer arrow shafts as having unequal lengths and think that my perceptual experience is not veridical, then I am not justified in believing in their inequality; and in fact I do not believe in that inequality if other perceptual experiences of the arrow shafts, experiences I deem veridical (such as the experience of their measurement, of superimposing one on the other, or of eliminating the arrowheads and flights),

show me that the arrow shafts are of equal length. And what holds for the Müller-Lyer illusion holds equally for the contents of all the perceptual experiences recognized by the perceiver as non-veridical. Recognizing the content of a perceptual experience as non-veridical automatically implies recognition of the unreal nature of its content; whereas recognizing a perceptual experience as veridical is equivalent to accepting that belief in the objectivity of its content is justified.

If in a trial counsel for the defence asks a witness how he can be so sure that the accused is the murderer, and the witness replies that he saw the accused commit the murder, then defence counsel can only argue against this justification by accusing the witness of lying or trying to show that his perception was non-veridical – for example, because he often suffers hallucinations, because he is a drug addict and was under the influence of drugs at the time of the crime, or because he has some visual defect that would have prevented him from correct identification of what he saw. What counsel cannot do is accept that the perception of the witness was veridical and not accept the justification based on that perception.


Our problem now would be to provide a criterion for distinguishing between veridical and non-veridical perceptual experiences. However, this issue has already been tackled in Chapter III of *Mente y mundo*. Briefly, it is argued there that the perceptual experience of any item is veridical if it is intersubjectively compatible with the remainder of the experiences of the same item. The experience of perceiving the Müller-Lyer arrow shafts as of unequal lengths is non-veridical because it is incompatible with the other perceptual experiences of the shafts. That is, perception sometimes deceives us, but we should not forget that it is also through perception that we recover from our mistakes.

It was also argued in Chapter III of *Mente y mundo* that in the case of non-veridical perceptual experiences we usually have an explanation for why we were led to a mistaken perceptual identification. This may be considered as a criterion that is subsidiary to the compatibility criterion.

There are two further points to which I wish to draw attention in relation to the distinction between veridical and non-veridical perceptual experiences. The first is that no perceptual experience is ultimate in the sense of being absolutely incorrigible: mistakes are always possible. Secondly, it is only *a posteriori*, by comparison with the rest of our perceptual experiences, that we can distinguish veridical from non-veridical perceptual experiences. While a perceptual experience is taking place, it is impossible to judge by its content whether it is veridical or non-veridical, except in special cases such as known illusions; and the same goes for hallucinations. The content of an hallucinatory perceptual experience, as such, can be indistinguishable from that of a veridical perceptual experience. What allows a veridical perceptual experience to be

distinguished from an hallucination is the intersubjective incompatibility of the latter with the rest of the perceptual experiences of those who share the same categorial frameworks; and, as I said, this can only be known *a posteriori* and by comparison.

After these clarifications of the proposed truth criterion, we can now see how veridical perceptual experiences can in fact justify the truth value of perceptual statements, and hence of the beliefs expressed by these statements. Statements are formed by concepts, and the concepts used in the perceptual statements of a given language *L* are the mental representation of the perceptual categories in which world *M* presents itself to the users of that language in the processes of perception. In other words, the concepts used in the constitution of the perceptual statements of a language *L* refer directly to perceptual categories of world *M*. The perceptual categories, together with the structural relationships among them, are their referents. The perceptual statement “The bird in the cage is a goldfinch” is true, and thus empirically justified, if there exist possible veridical perceptual experiences such that the statement in question expresses, at the symbolic level, the categorial and structural contents of the corresponding perceptual categories. The concepts and structural relationships expressed by the statement refer correctly to the corresponding perceptual categories. The statement “The bird in the cage is a goldfinch” is true if at the perceptual level it is in fact verifiable that there is indeed a bird in the cage and that this bird is a goldfinch. Once this has been verified, and taking for granted that this verification is the result of a veridical perceptual experience, we can conclude that the belief expressed by the statement is true and empirically justified.



VIII. Categorical perception and scepticism



If perceptual statements are empirically “justified” by perception, as was argued in the previous chapter, it seems logical for me to provide a response to philosophical scepticism regarding perception, to those who hold that through perception we can know nothing of the world around us. What motivates this kind of scepticism is the sceptic’s belief that we do not access the world directly through perception, but only via our representations of the world. Even in perceptual experiences occurring under the most favourable conditions, as when we see a tomato on a table or a bird in a cage, there is always a chance that we are experiencing an hallucination, or that what we seem to be seeing is not really a tomato or a bird; and according to the sceptic this is possible because we do not access the world directly in perception, but only via its representation, a representation that, as such, for the perceiving subject can be exactly the same for the perception of something real as for the perception of something unreal.

The sceptic’s conclusion makes some kind of sense, because his view of perception presupposes that the objective world is in itself one thing, and our representations of it in perception are another. It can always happen, as in illusions, hallucinations and mistaken perceptual identifications, that our perceptual representations do not correspond to what is really there; since according to the sceptic we have immediate access only to the world of our representations and not to real world represented by them, he concludes that we have no means of knowing whether these representations do or do not correspond to how things are in the real world. In short, the sceptic thinks that we cannot know whether everything we experience perceptually as real, as existing out there, independently of us, is actually no more than a dream of reality; he sees no bridge across the divide that separates our representations from what lies beyond the mental world. In the words of John McDowell (2006), “The familiar sceptical scenarios - Descartes’s demon, the scientist with our brains in his vat, the suggestion that all our apparent experience might be a dream - are only ways to make this supposed predicament vivid”; the predicament being how to cross the divide that separates our supposed perceptual representations from that of which they are or purport to be representations.

If we accept the dualist premise that underlies sceptical arguments, I would say there is indeed no escape from scepticism, however wrong and contrary to common sense it may seem. From the moment we grant that the contents of perception are representational, making representation a kind of intermediary between us and the world, the representation - and with it we ourselves - will always be on this side of the divide. Only the Divine Eye can see both sides, and

consequently only He can *know* whether the representational content of one side do or do not correspond to the real world of the other side. We human beings know our representations, but according to the sceptic we do not *know, sensu stricto*, anything at all about the objective world around us - not even whether it exists.

The forcefulness of the sceptical argument stems from the supposition that there is a possible world beyond all our perceptual experiences. According to the sceptic, perception does not afford us immediate access to the real world, but to our representation of it; and there is evidently always a chance that this representation does not correspond to what is really out there, existing independently of us. It is even possible that that supposed real world does not exist, as occurs in hallucinations. Why can it not be that everything we experience perceptually as real has no more reality than what it is endowed with by our imagination? These sceptical doubts, however absurd they may seem, indicate something plausible; but its evident plausibility rests on two premises: firstly, the existence of an independent world (in spite of this existence possibly being questioned in the course of the sceptic's subsequent reasoning); and secondly, that in the processes of perception we do not gain immediate access to that supposedly independent world, but to the representations of it that we humans make for ourselves. If we accept these two premises, the sceptic's arguments, leading to the conclusion that we can have no certainty about the validity of our knowledge of the world, are as far as I can see irrefutable.

But what grounds does the sceptic have for accepting the validity of those two premises? Let us start with the first, the assumed existence of an independent world. The independent world of the sceptic cannot be the world we access through our senses, and which we describe as external to and independent of us, since according to the sceptic this world is the world of our representations. And in this the sceptic is evidently right, it is the world as it is experienced by cognitive beings in the processes of perception or experimental scientific observation; but it is precisely of this world that we say, and of which it makes sense to say, that it is external and independent. What grounds does the sceptic have for supposing that the *really* real world, if it exists, is not this one? His only grounds can only be the supposition - his second premise - that perception does not afford us immediate access to the real world of things that are independent of mind, but to the representations that we make of them. Here is where the sceptic commits his second error; but for the moment let us go on examining the first.

To suppose or postulate, as the sceptic does, the existence of an independent world "in itself", different from the world we experience perceptually as independent, is not justified in the sceptic's discourse. And it cannot validly be replied that the sceptic does not assume the existence of any independent world, but in fact questions that existence, claiming that it could well be that all the

external world of which we become conscious in the processes of perception is no more than a figment of our imagination. The sceptic does indeed question not only the validity of our representations of the world around us, but also the very existence of the world represented by the contents of our perceptual experiences; but to initiate his sceptical argument he must start from the premise of the possible existence of a real world distinct from the experienced world, even though in the course of his reasoning he later questions that existence. To get his argument going he must start by assuming, albeit provisionally, the existence of an independent “world in itself” that is distinct from the perceptually experienced world. And I ask: “What grounds does he or could he have for accepting this premise?”

One reply, very much in the Western tradition exemplified by the scepticism of Descartes, consists in postulating that there is indeed an external world created by God. In this context, the problem for we humans is to know whether the world we know corresponds to the world designed by its creator; in other words, whether the world as known with absolute certainty by the God who created it matches the world as human cognitive beings believe it to be on the basis of what we have learnt of it through perception or scientific observation. In this context arguments in favour of scepticism become meaningful because the nature of the independent world is well defined as being what its creating cognitive being endowed it with. Given this, it is always possible that the knowledge of that world that is possessed by other cognitive beings, in this case human beings, does not conform to what it really is, to what it really is for the cognitive being that created it.

In the absence of this context, the premise of a world endowed with a structure and properties that are independent of how they are experienced by cognitive beings is utterly baseless. This is not to say that the structure and properties of the world that are manifested to cognitive beings do not belong to that world independently of the particular experience that a particular cognitive being may have of them; but here, “independently” merely means that those properties and structure, which unlike mental contents belong to things and not to our representation of them, belong to things *qua* things that can be experienced by cognitive beings. To speak of an assumed reality of the world that is independent of how the world manifests itself to cognitive beings is meaningless, it is to speak of a chimera, and the sceptic is accordingly right in asserting that we can know nothing for certain about that chimerical, supposedly independent world. But the consequences of postulating the existence of such a world do not stop there, since it is this very postulate that forces the sceptic to consider the contents of perception as representational.

Only in so far as we take for granted the existence of a “world in itself”, endowed with a reality and existence different from the reality and existence

with which the world presents itself to cognitive beings in the processes of perception, are we obliged to think of the content of our perceptual experiences as representational, because the content of perceptual experience is one thing and the independent world represented by that experience is another. Once we have assumed this kind of Cartesian dualism, in which mind and world occupy parallel spaces, it is obvious that all we can find in mind about the world is its representation; and as the sceptic argues, there is no way of knowing whether this representation matches what it represents. It is even possible that outside the mind the world does not even exist.

The sceptic's reasoning is flawless; but as I said before, it all rests on two highly questionable premises, an independent "world in itself" and the supposedly representational character of perceptual contents. Now we are not going to deny the existence of the outside world here. What seem baseless and untenable are the conception of the outside world that the Cartesian dualist has wrought for himself, and the consequent parallel conception of perceptual contents as representational.

If we shake off the dualist trammels that shackle the sceptic, forcing him to regard perceptual content as representational, we find ourselves in a quite different situation. To speak of representation implies, at the least, the existence of two things: the representation and what is represented. That is why, at the level of mental contents, I consider the content of beliefs and memories to be representational. The content of a belief or memory is one thing, and what it is a belief or memory of is another. I may believe that it will rain tomorrow because that is what the weatherman said, but the content of my belief is one thing, and that tomorrow it will or will not rain is another. And the same goes for memories, *mutatis mutandis*: I may think that I connected the alarm on leaving home this morning, but the content of my memory is one thing, and that I did or did not connect the alarm is another. But in both cases, beliefs and memories, both kinds of content are accessible to the subject, and that is why we can consider one of them to be a representation of the other. Trusting the weatherman I believe it will rain tomorrow, and tomorrow I will be able to check whether the content of my belief matches the event represented or not.

If, free of the dualist trammels that shackle the sceptic, we now analyse what goes on in perception, we find something quite different from the cases of beliefs and memories. Whereas the mental contents of beliefs and memories manifested themselves to consciousness as representing something different from themselves, with perception it is not so. When I have the perceptual experience of seeing something, such as the keyboard at which I am writing, I am only aware of one thing, the thing I am seeing, which I experience as being out there before me with its own existence and without at all depending on me to keep on existing. The content of my perceptual experience of the keyboard is nothing

else but the way the keyboard has of presenting itself to my consciousness in the process of perception. If I tried to describe the content of my perceptual experience, what I would really be describing would be the keyboard, not a representation of it. And this presentational content of the keyboard, unlike the content of beliefs and memories, does not refer to anything other than itself and therefore cannot be said to be representational. Only in relation to the chimerical world of the sceptic does it acquire that supposedly representational character; but that is not the way it presents itself to cognitive beings in the processes of perception. In short, as I have already argued in Chapter VI, in perception things present themselves to us, we have direct access to them, understanding by this that it is not direct access to the chimerical world of the sceptic that we enjoy, but direct access to the world as it is processed by the architecture of our perceptual apparatus; and it is of the world so processed that it is possible to say meaningfully that it is real, that it is out there before us, and that its existence does not depend on what we may or may not perceive.

As I said, in the precise moment in which I am perceiving the keyboard of my computer, there is only one thing present in my consciousness, the keyboard, and it is the *memory* of this presentational content that is representational. And what does it represent? Obviously, the keyboard, exactly as it presents itself in the corresponding acts of perception. It is thus that the word “representation” makes sense, when we have at our disposal both members of the representational relationship. First we perceptually identify things in the world, and the recollection of the categorial content of these presentations is representational. We ask again: “Representational of what?” Of the presentational content of perception.


Unlike the sceptic, who uses the term “representation” for something, the referent of which he denies having access to, here we have at our disposal both members of the relationship of representation: on the one hand, the mental content of the representation; and on the other, what is represented, the items of the world that have been identified perceptually. We are of course thinking here of the initial, most basic representations of the world: later, with the aid of imagination, we can construct representations of fictitious entities, or of entities we suppose to be real but are not accessible to the senses.

Since our access to the world is direct, since in perception the things of the world present themselves to us, there is no call for scepticism. Scepticism consists in the dualistic premise that mind and world occupy two different spaces with no mutual communication, and a conception of knowledge as an attempt to bridge the abyss that separates those spaces. Knowledge does indeed imply the existence of two different spaces: the space of mental representations, and the space of the world. But this world, as the present analyses of perception have shown, is not the chimerical world of the sceptic but the world of our

experience, the world as it is processed by the architecture of our perceptual system. What the world may be, other than as it can or does present itself to cognitive beings in the processes of perception or experimental scientific observation, is a meaningless question, because it has no possible answer.

Of course, the world is not just what is perceived or observed. It manifests itself, in the processes of perception themselves, as transcendent, as not limited to what we may know of it at any given time. But in the last instance it is always the world as experienced that decides the truth value of our representations, not the chimerical world of the sceptic.

And finally, let us remember that it is of the world of our veridical perceptual experiences and scientific observations that we predicate existence and reality, with the caveat that, as discussed in previous chapters, no perceptual experience or observation is ultimate in the sense of being incorrigible. So if what we want to understand by knowledge of the world is ultimate and definitive knowledge, we are thinking of the sceptic's chimerical "world in itself", and are asking for something that, effectively by our own definition, is unattainable. That is not the world of cognitive beings such as ourselves.



Bibliography

- Barsalou, L. W. (1983): "Ad hoc categories" in *Memory and Cognition*, 11, 211-227.
- Barsalou, L. W. (1987): "The instability of graded structure: Implications for the nature of concepts", in U. Neisser (ed.): *Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization*, Cambridge, Cambridge University Press.
- Berlin, B. (1992): *Ethnobiological Classification: Principles of Categorization of Plants and Animals in Traditional Societies*, Princeton, NJ, Princeton University Press.
- Berlin, B., Breedlove, D. E., & Raven, P. H. (1973): "General principles of classification and nomenclature in folk biology", *American Anthropologist*, 75, 214-242.
- Berlin, B. and Kay, P (1969): *Basic Color Terms: Their Universality and Evolution*, Berkeley CA, University of California Press.
- Carlson, Neil R. (1999): *Fisiología de la conducta*, Barcelona, Ed. Ariel.
- Cohen, B. and Murphy, G. L. (1984): "Models of concepts", *Cognitive Science*, 8, 27-58.
- Davidson, D. (1989): "A Coherence Theory of Truth and Knowledge" in Ernest LePore (ed.) *Truth and Interpretation. Perspectives on the Philosophy of Donald Davidson*, Oxford, Basic Blackwell, pp. 307-319.
- Fuster, J. (2003): *Cortex and Mind. Unifying Cognition*, Oxford, New York, Oxford University Press.

- Hampton, J. A. (1979): "Polymorphous concepts in semantic memory", *Journal of Verbal Learning and Verbal Behaviour*, 18, 441-461.
- Humphreys, G. W & Riddoch, M. J. (1987): "To see but not to see: a case study of visual agnosia", London, Lawrence Erlbaum Associates.
- Kanizsa, G. (1980): *Gramática del vedere. Saggi su percezione e gestalt*, Società Editrice Il Mulino, Bologna (Spanish edition.: *Gramática de la visión. Percepción y pensamiento*, Paidós, Barcelona, 1986).
- Lamberts, K. and Shanks, D. (1997): *Knowledge, Concepts and Categories*. London, UCL Press.
- LePore, E. (ed.) (1989): *Truth and Interpretation. Perspectives on the Philosophy of Donald Davidson*, Oxford, Basic Blackwell.
- Malt, B. C. (1995): "Category coherence in cross-cultural perspective", *Cognitive Psychology*, 29, 85-148.
- Mandler, J. M. & McDonough, L. (2000): "Studies in inductive inference in infancy", *Cognitive Psychology*, 37, 60-96.
- Markman, Ellen M. (1989): *Categorization and Naming in Children: Problems of Induction*, Cambridge, MA, MIT Press.
- McDowell, John (2006): "The Disjunctive Conception of Experience as Material for a Transcendental Argument", *Teorema*, vol. XXV/I, pp.19-33 (Spanish translation: "La concepción disyuntiva de la experiencia: material para un argumento transcendental", *Cuadernos de pensamiento*, 2, 2007, pp. 83-121).
- Medin, D. L. and Schaffer, M. M. (1978): "Context theory of classification learning", *Psychological Review*, 85, 207-238.

- Metzinger, Thomas (2003): *Being No One. The Self-Model Theory of Subjectivity*, Cambridge, Massachusetts, The MIT Press.
- Murphy, Gregory L. (2004): *The big book of concepts*. Cambridge, Massachusetts, MIT Press.
- Murphy, G. L. and Smith, E. E. (1982): "Basic level superiority in picture categorization", *Journal of Verbal Learning and Verbal Behavior*, 21, 1-20.
- Putnam, H. (1981): *Reason, Truth and History*, Cambridge, Cambridge University Press (Spanish edition: *Razón, verdad e historia*, Ed. Tecnos, Madrid, 1988).
- Putnam, H. (1990): *Realism with a Human Face*, Cambridge, Harvard University Press.
- Prinz, Jesse, J. (2004): *Furnishing the mind. Concepts and their perceptual basis*, Cambridge, Massachusetts, The MIT Press.
- Raffman, Diana (1995): "On the persistence of phenomenology", in Metzinger, T.: *Conscious Experience*, Schöningh / Imprint Academic, 293-308.
- Rosch, Eleanor (1975): "Cognitive representation of semantic categories", *Journal of Experimental Psychology: General*, 104, 192-233.
- Rosch, E. and Mervis, C. B. (1975): "Family resemblance: Studies in the internal structure of categories", *Cognitive Psychology*, 7, 573-605.
- Rosch, E., Mervis, C. B., Gray, W., Johson, D., and Boyes-Braem, P. (1976): "Basic objects in natural categories", *Cognitive Psychology*, 8, 382-439.
- Rosch, Eleanor and Lloyd, Barbara B. Eds.) (1978): *Cognition and Categorization*, Lawrence Erlbaum Associates, Publishers, Hillsdale, New Jersey.
- Smith, Edward (2008): *Procesos cognitivos: modelos y bases neurales*, Madrid, Pearson.

- Smith, E. E. and Medin, D. L. (1981): *Categories and Concepts*, Cambridge, MA, Harvard University Press.
- Tanaka, J. W. & Curran, T. (2001): "A neural basis for expert object recognition", *Psychological Science*, 12, 43-47.
- Vázquez, J. (1998): "How language hooks on to the world", in Martínez, C., Rivas, U. And Villegas, L. (Eds.) (1998): *Truth in Perspective. Recent Issues in Logic, Representation and Ontology*, Aldershot, Ashgate Publishing Limited.
- Vázquez, J. (2000): "Semántica de los nombres propios, deícticos y términos de clase", *Teorema*, Vol. XIX/1, pp. 75-92.
- Vázquez, J. (2000): "Epistemic Truth in a Plurality of Worlds", *Logica Trianguli*, 4, pp.53-67.
- Vázquez, J. (2004): "La observación científica en el proceso de contrastación de hipótesis y teorías", *Theoria*, 49, pp. 77-95.
- Vázquez, J. (2007): *Mente y mundo. Aproximación Neurológica*, Madrid, Akal.
- Vázquez, J. (2010): "Los enunciados de percepción y su valor de verdad (La justificación empírica del conocimiento)", in Sofia Miguens, Joao Alberto Pinto, Manuela Teles (eds.): *Aspectos do Juízo (Aspects of Judgement)*, U. Porto, Porto, pp.135-162.

Putting forward an original analysis of perceiving as a cognitive attitude, as it contrasts with judging, believing and knowing, the author approaches several issues in the philosophy of perception, such as differences between presentation and representation, the natures of concepts and categorization, the justification of perceptual beliefs and their role in the justification of knowledge. Influenced by phenomenology and by psychology and neuroscience of vision, Juan Vázquez's view brought to Project *Bounds of Judgement* a standpoint from which to discuss perceptual disjunctivism, our reference position for work in the field.

Sofia Miguens

The present volume is a joint outcome of Portuguese and Spanish projects *The Bounds of Judgement – Frege, cognitive agents and human thinkers* (PTDC/FIL-FIL/109882/2009) and *Cognitive Attitudes and the Justification of Knowledge* (FFI2009-08828) (Principal Investigators: Sofia Miguens – University of Porto and Concha Martínez – University of Santiago de Compostela).

Juan Vázquez has been Professor of Logic and Philosophy of Science (currently ad honorem Professor) at the Department of Logic and Moral Philosophy of the University of Santiago de Compostela (Spain). He is the author of *Aproximación analítica al pensamiento platónico*, 1982; *Lenguaje, verdad y mundo*, 1986; *Mente y Mundo. Aproximación neurológica*, 2007; as well as a large number of articles on philosophy of science, epistemology and philosophy of language. He has been collaborating with the Mind, Language and Action Group (MLAG) of the Institute of Philosophy of the University of Porto since its beginning.