COGNITIVE SCIENCE

Cognitive Science 25 (2001) 423–451

http://www.elsevier.com/locate/cogsci

Is perception informationally encapsulated? The issue of the theory-ladenness of perception

Athanassios Raftopoulos*

Assistant Professor of Philosophy and Cognitive Science, Department of Educational Sciences, University of Cyprus, P.O. Box 20537, Nicosia 1678, Cyprus

Abstract

Fodor has argued that observation is theory neutral, since the perceptual systems are modular, that is, they are domain-specific, encapsulated, mandatory, fast, hard-wired in the organism, and have a fixed neural architecture. Churchland attacks the theoretical neutrality of observation on the grounds that (a) the abundant top-down pathways in the brain suggest the cognitive penetration of perception and (b) perceptual learning can change in the wiring of the perceptual systems. In this paper I introduce a distinction between sensation, perception, and observation and I argue that although Churchland is right that observation involves top-down processes, there is also a substantial amount of information in perception which is theory-neutral. I argue that perceptual learning does not threaten the cognitive impenetrability of perception, and that the neuropsychological research does not provide evidence in favor of the top-down character of perception. Finally, I discuss the possibility of an off-line cognitive penetrability of perception. © 2001 Cognitive Science Society, Inc. All rights reserved.

1. Introduction

In his article "Observation reconsidered" Fodor (1984, 120)[1] argued that observation is theory neutral, since "two organisms with the same sensory/perceptual psychology will quite generally observe the same things, and hence arrive at the same observational beliefs, however much their theoretical commitments may differ" (emphasis in the text).

Later in his article Fodor (1984, 127) concedes that a background theory is inherent in the process of perceptual analysis. In this sense, observation is theory-laden, and inference-like. But, this theory-ladenness does not imply that observation is theory-dependent in the way relativistic theories of philosophy of science and holistic theories of meaning intended it to

---

* Tel.: +357-2-753749.

E-mail address: raftop@ucy.ac.cy (A. Raftopoulos).

PII: S0364-0213(01)00042-8

be. The reason is that these theories require that the perceptual analysis have access to background knowledge, and not just to the theory that is inherent in the system. But this is not true in view of the various implasticities of perception (as the Muller-Lyer illusion), which show that how things look is not affected by what one believes. This argument is best understood in the light of Fodor's (Fodor, 1983) view regarding the modularity of the perceptual systems, that, unlike reflexes, they are computational but informationally encapsulated from information residing in the central neural system.

The input systems, or perceptual modules, are domain-specific, encapsulated, mandatory, fast, hard-wired in the organism, and have a fixed neural architecture. Their task is to transform the proximal stimuli that are registered by the sense organs to a form that is cognitively useful, and can be processed by the cognitive functions. This transformation is a computation which relies on some general assumptions, whose role is to reduce the sensory ambiguity and allow the extraction of unambiguous information from the visual array.

The perceptual modules have access only to background theories that are inherent in these modules. The modules, in this view, do not have access to our mental states and theoretical beliefs. Fodor (1984, 135) distinguishes between "fixation of appearances" or "observation," which is the result of the functioning of the perceptual modules, and "fixation of belief," which is the result of the processing of the output of the modules from the higher cognitive centers. The former is theory-encapsulated, the latter is not. Fodor's target was the *New Look* theories of perception, according to which there are no significant discontinuities between perceptual and conceptual processes. Beliefs inform perception as much as they are informed by it, and perception is as plastic as belief formation.

Fodor's (Fodor, 1984) argument is that, although perception has access to these background theories and is a kind of inference, it is impregnable to (informationally encapsulated from) higher cognitive states, such as desires, beliefs, expectations, and so forth. Since relativistic theories of knowledge and holistic theories of meaning argue for the dependence of perception on these higher states, Fodor thinks that his arguments undermine these theories, while allowing the inferential and computational role of perception and its theory-ladenness.

Churchland (1988)[2] attacks Fodor's views about theoretical neutrality of observation on two grounds. He argues, first, that perceptual implasticity is not the rule, but rather the exception, in a very plastic brain, in which there is ample evidence that the cognitive states significantly affect perception. Thus, he rejects the modularity of the perceptual systems. Second, he claims that even if there is some rigidity and theoretical neutrality at an early perceptual process, this "pure given," or sensation, is useless in that it cannot be used for any "discursive judgment," since sensations are not truth-valuable, or semantically-contentful states. Only "observation judgments" can do that, because they have content, which is a function of a conceptual framework. Thus, they are theory-laden.

In this paper I will address Churchland's claim about the plasticity of the perceptual systems, his arguments against their modularity, and assess their effectiveness as a critique of the theoretical neutrality of observation. The conclusion is that though Churchland is right that observation involves top-down processes, there is also a substantial amount of information in perception which is clearly bottom-up, and theory-neutral. I shall not argue here whether this theory-neutral perceptual basis is semantically contentful, and whether it can be used as a theory-neutral basis for assessing competitive theories.

The problem with both Fodor and Churchland is their conception of the sensation-perception-cognition distinction, and an adequate account of this distinction will help us delineate what exactly is at stake in their arguments. Both approaches, moreover, view perceptual learning and the structural changes it induces, as a threat to the cognitive impenetrability of the modules. Fodor, because he thinks that the input systems have a fixed neural architecture, and Churchland, because he thinks that perceptual learning demonstrates the cognitive penetrability of perception. Both views are wrong.

Finally, Churchland's claims that recent neuropsychological research provide evidence in favor of the top-down character of perception will be addressed. I will claim that Churchland misinterprets this evidence and that these findings can be reconciled with a modularized view of human perceptual systems.

Before proceeding some terminological discussion is in order. The terms "perception" and "observation" will frequently be employed in this paper and will be carefully distinguished the one from the other. These terms are not employed consistently in the literature. Sometimes "perception" purports to signify our phenomenological experience, and thus, "is seen as subserving the recognition and identification of objects and events" (Goodale, 1995, 175). In Goodale's sense, "perception" is a wider process, which includes "observation." Since these terms are not used in the same way in this paper—I somewhat modify Shrager's (1990, 439) usage, according to which "perception" connects "sensation" to "cognition"—some terminology will be introduced now with a view to explicating the vocabulary used in this paper.

The term "sensation" refers to all processes that lead to the formation of the retinal image (the retina's photoreceptors register about 120 million pointwise measurements of light intensity). This set of measurements, which initially is cognitively useless, is gradually transformed along the visual pathways in increasingly structured representations that are more convenient for subsequent processing. The processes that transform sensation to a representation that can be processed by cognition are called perception. Perception includes both low-level and intermediate-level vision, and is bottom-up, in that it extracts information retrievable directly from the scene only. In Marr's (Marr, 1982) model of vision, which is discussed below, the *21/2D* sketch is the final product of perception. In other models, perception encompasses all processes following sensation that produce their output independent (in an on-line manner) of any top-down information, although information from higher levels may occasionally select the appropriate output. All subsequent processes are called cognition, and include both the postsensory/semantic interface at which the object recognition units intervene as well as purely semantic processes. At this level we have observation. The formation of Marr's *3D* model, for example, is a cognitive activity.

2. Are perceptual systems informationally encapsulated?

2.1. The argument from illusions

Churchland's first argument against the impenetrability thesis of the perceptual systems consists in an examination of various illusions and visual effects, such as the Necker cube, the well-known rabbit-duck figure, and so forth, which reveal that there is "a wide range of elements central to visual perception–contour, contrast, color, orientation, distance, size, shape, figure versus ground–fall of which are cognitively penetrable" (Churchland 1989, 261).

The interpretation of visual illusions is controversial. There is disagreement about their causes and the extent to which their resolution depends on top-down flow of information. The whole issue hinges on exactly how much of a top-down process vision is, and what is the nature of the top-down influences. In what follows I will use Marr's (Marr, 1982) theory of vision as an example of the kind of modular theory that Fodor is arguing for, to show how Churchland's observations concerning illusions can in fact be accommodated in a semi-Fodorian framework.

2.1.1. Top-down and bottom-up processes in vision

According to Marr (1982), there are three levels of representation. The initial level of representation is the primal sketch, which captures contours and textures in an image. The second level is the observer-centered $2^{1/2}D$ sketch, which yields occlusion relations, orientation and depth of visible surfaces. Recognition of objects requires an object-centered representation, Marr's *3D* model.

Marr considers the $2^{1/2}D$ sketch to be the final product of the bottom-up, data- driven early vision, that is, perception. Its aim is to recover and describe the surfaces present in a scene. Visual processes that process information such as surface shading, texture, color, binocular stereopsis, and analysis of movement are referred to as low-level vision. Its stages purport to capture information that is extractable directly from the initial optical array without recourse to higher level knowledge.

Hildreth and Ulmann (1989) argue for the existence of an intermediate level of vision. At this level occur processes (such as the extraction of shape and of spatial relations) that cannot be purely bottom-up, but which do not require information from higher cognitive states. These tasks do not require recognition of objects, but require the spatial analysis of shape and spatial relations among objects. This analysis is task-dependent, in that the processes involved may vary depending on the task being accomplished, even when the same visual array is being viewed.

The recovery of the objects present in a scene cannot be the result of low-level and intermediate-level vision. This recovery cannot be purely data-driven, since what is regarded as an object depends on the subsequent usage of the information, and thus is task-dependent and cognitively penetrable. In addition, most computational theories of vision (Marr, 1982; Biederman, 1987) hold that object recognition is based on part decomposition, which is the first stage in forming a structural description of an object. It is doubtful, however, whether this decomposition can be determined by general principles reflecting the structure of the world alone, since the process appears to depend upon knowledge of specific objects (Johnston, 1992). Object recognition, which is a top-down process, and requires knowledge about specific objects is accomplished by the high-end vision. Object recognition requires a matching between the internal representation of an object stored in memory and the representation of an object generated from the image. In Marr's model of object recognition the

3D model provides the representation extracted from the image which will be matched against the stored structural descriptions of objects (perceptual classification).[3]

Against Marr's model of object recognition, Lawson, Humphreys, and Watson (1994) argue that object recognition may be more image-based than based on object-centered representations, which means that the latter may be less important than Marr thought them to be. Neurophysiological studies by Perrett et al. (1994) also suggest that both object-centered and viewer-centered representations play a substantial role in object recognition.

Other criticisms address the issue of Marr's thesis regarding functional modularity, that is, the idea that a large computation can be split up and implemented as a collection of parts that are as nearly independent of one another. There is ample evidence (Cavanagh, 1988; Gilchrist, 1977; Livingstone & Hubel, 1987) that the early vision module consists of a set of interconnected processes (submodules) for shape, color, motion, stereo, and luminance that cooperate within it. These are functionally independent and process stimuli in parallel. Thus, early vision consists of a continuum of multiple, parallel, task specific, modules. This, internal to early vision, "horizontal" or "lateral" flow of information, however, does not threaten the cognitive impenetrability of early vision, since it leaves no room for penetration of knowledge from higher extarvisual cognitive centers.

Neurophysiological evidence (Felleman et al., 1997) also suggests that information flows in a top-down manner from loci higher along early vision to earlier stages of early vision. Being within the early vision module, however, this top-down flow of information does not support the cognitive penetrability of early vision from extravisual information.

Thus, despite criticisms of Marr's program, his distinction between early representations, that are most likely bottom up, and higher level representations, that are informed by specific knowledge, remains valid. His notion of functional modularity also holds, provided that one views Marr's modules as consisting of a set of submodules with lateral and top-down channels of communication that process in parallel different information extracted from the retinal image.

There is a host of neurological and neuropsychological findings that support the above conclusion. Consider the cases of visual object agnosia. Visual agnosias can occur for different kinds of stimuli (colors, objects, faces), and may affect either the ability to copy or the ability to recognize objects. Research (Newcombe & Ratcliff, 1977; Warrington & Taylor, 1978; Warrington, 1975, 1982; Humphreys & Riddoch, 1984, 1985; Campion & Latto, 1985) shows that there is a relative autonomy of the components of the visual processing routines. Damage to the early visual routines causes impairments at high-level vision, but damage to high-level vision usually leaves low vision intact.

Impairments of the object-centered representation, for instance, leave intact the lower viewer-centered representation. More specifically, difficulty in identifying objects that are seen from unusual views (Warrington & Taylor, 1973), difficulties in matching by physical identity (Warrington & Taylor, 1978), as well as difficulties in recognizing that an object has the same structure even when its view changes–object constancy (Humphreys & Riddoch, 1984, 1985), suggest an impairment in the formation of the *3D model* (the object-centered representation). In so far as the patients perform normally in categorization tasks, their viewer-centered representation is intact.

Semantic memory impairments leave intact both the initial, viewer-centered and object-centered representation. Damage in the left hemisphere (De Renzi, Scotti & Spinnler, 1969) is accompanied by the so-called semantic impairments, in which knowledge of the objects' category, classification, of properties and functions is degraded or inaccessible. Studies (Taylor & Warrington, 1971; Warrington, 1975) show that the same patients have normal initial, viewer-centered, and object-centered representations, since they succeed in matching tasks, drawing objects, recognizing objects seen from unusual views, and maintain object constancy. Thus, the semantic impairments affect neither perception, nor observation (the formation of the object-centered representation).

The neuropsychological evidence provided by studying the various forms of visual agnosias suggests that there is a relative autonomy of the various components of the visual processing routines. Damage to the early visual routines causes impairments at high-level vision, but damage to the high-end routines usually leaves low vision intact. Thus, although sensory visual deficiencies and/or impairment in the formation of the *primal sketch* affect performance in tasks that require the $2^{1/2}$ *sketch* and/or *3D model,* impairments at this latter level are not reflected in deficiencies in tasks depending on the formation of the *primal sketch.* Similarly, impairments of the object-centered representation leave intact the lower viewer-centered representation. Finally, semantic memory impairments that affect object categorization and deprive the patient access to her background object-specific knowledge seem to leave intact not only the *perceptual* initial and viewer-centered representations, but also the more cognitive object-centered representation. Therefore, the various parts of the visual processing system seem to satisfy Marr's principle of modular organization, and they seem to enjoy a cognitive impenetrability *ala* Fodor.

Before proceeding a remark is in order. Most programs of vision are concerned with visual processes that culminate in high-level vision, by means of which visual patterns are recognized, objects are identified and categorized. But, in evolutionary terms, vision's most important function is to guide action.

Now, two prominent pathways, or streams of visual projections, have been identified in the primate cerebral cortex. The one is the ventral stream, which stems from the primary visual cortex and projects to the inferotemporal cortex, and the other is the dorsal stream, which stems from the visual cortex and projects to the posterior parietal cortex. Current research (Goodale, 1993, 1995; Goodale & Milner, 1992) suggests that these two distinct streams play different roles in vision, although the two systems collaborate to produce our complex behavioral patterns. The ventral stream is responsible for vision in perception and the dorsal stream for vision in action. More specifically, the ventral stream leads to the formation of perceptual representations, that is, the percepts on the basis of which we build our knowledge about the world. The dorsal stream mediates the visual control of our skilled actions.

The same research also indicates that the information processing taking place along the dorsal pathway is very likely impervious to top-down flow of information (it is nonconceptual). Since the controversy regarding the cognitive penetrability of vision concerns the extent to which the percepts are informed from a top-down flow of information, I will restrict my discussion to the visual path that leads to the formation of the percepts. Thus, the discussion will focus on the ventral visual pathway, that is on vision in perception.

2.1.2. Implicit versus explicit knowledge and cognitive penetrability

According to Fodor (1983) perception has access to some background theories and is a kind of inference. Still, it is impregnable to higher cognitive states, such as desires, beliefs, expectation. A distinction is drawn, thus, between the theories informing perception and the specific knowledge about objects that constitutes the representational content of our cognitive states (beliefs, expectations, desires).

The computations involved in all levels of vision are constrained by some principles. These constraints are needed, because perception is underdetermined by any particular retinal image; the same retinal image could lead to distinct perceptions. The problem is accentuated with regard to the underdetermination of the $2^{1/2}D$ structure (three dimensional) from the *2D* retinal stimulation (two-dimensional). Unless the observer makes some assumptions about the physical world which gives rise to the retinal image, perception is not feasible. Thus, even if perception is bottom-up, still it is not insulated from knowledge. Knowledge intrudes on perception, since early vision is informed and constrained by some general world principles that reduce indeterminancies in information.

Among these principles are those of "local proximity" (adjacent elements are combined), of "closure" (two edge-segments could be joined even though their contrasts differ because of illumination effects), of "continuity" (the shapes of natural objects tend to vary smoothly and usually do not have abrupt discontinuities), "compatibility" (a pair of image elements are matched together if they are physically similar, since they originate from the same point of the surface of an object), and "figural continuity" (figural relationships are used to eliminate most alternative candidate matches between the two images).

Most computational accounts (Ulmann, 1979; Marr, 1982) hold that these principles substantiate some general truths of our world and are not assumptions about specific objects acquired through experience. In this sense, they are general theories about our world. Moreover, they seem to be hardwired into the system. Thus, even the early stages of vision, in so far as they involve some built-in physical constraints, or theories, are theory-laden. These constraints provide the body of background knowledge stored in Fodor's perceptual modules. In this sense, and if one metaphorically interprets the processes involving the general constraints as "thinking," one could agree with (Spelke, 1988, 458) that "perceiving objects may be more akin to thinking about the physical world than to sensing the immediate environment." These principles however are not the result of explicit knowledge acquisition about specific objects but are general reliable regularities about the optico-spatial properties of our world hardwired in our perceptual systems.

This knowledge is implicit, in that it is available only for the processing of the retinal image, whereas explicit knowledge is available for a wide range of cognitive applications. Implicit knowledge cannot be overridden. The general constraints hardwired in the visual system can be overridden only by other similar general constraints with which they happen to compete (although no one knows yet how the system "decides" which constraint to apply). Still, one cannot decide to substitute it with another body of knowledge, even if one knows that under certain conditions this implicit knowledge may lead to errors (as is the case with the visual illusions). This theoretical ladenness, therefore, cannot be used as an argument against the existence of a theory-neutral ground, because perception based on a shared theory is common ground.

The physical constraints at work in perception must be reflected in the physiological mechanisms underlying the early stages of vision, since it is these mechanisms that implement them (Hildreth & Ulmann, 1989). There is evidence that the constraints applied to restrict the possible alternative solutions to computational problems of vision are reflected in the physiological mechanisms underlying binocular computations, from cells for edge detection to mechanisms implementing the epipolar constraint (Hubel & Wiesel, 1968; Poggio & Talbot, 1981; Ferster, 1981; Watt & Morgan, 1984; Koch & Poggio, 1987).

2.1.3. Illusions

Let me now return to the issue of illusions. These are cases in which the internal organization of a figure or a scene is such that it deceives one in constructing the wrong organization, or is so ambiguous that its resolution seems to rely on higher cognitive factors. Similar problems occur also in linguistics, when syntactic analysis yields, for instance, more than one possibilities for the grammatical role of a word.

Fodor (1983,76) is aware that the context may be invoked to solve certain syntactic ambiguities. The problem he considers is whether this top-down process determines the output of the syntactic module. If the top down flow of information affects the module prior to its production of output, then there would be no cognitive impenetrability. But Fodor believes that though the context solves the disambiguity, it does not determine the output of the modules. What happens is that the module proposes all possible syntactic analyses, and the higher-level processes determine which interpretation is acceptable. All others are deactivated and do not participate in further processing. Thus, “the effects of sentence context . . . must be . . . ‘postperceptual’ . . . these processes must operate after the input system has provided a (tentative) analysis of the lexical content of the stimulus.”

This “filtering interactive” model allows a weak interaction between higher and lower levels of processing, since the channel of communication between the levels allows a very limited feedback, the only information that passes is “yes,” that is, acceptable, or “no,” that is, unacceptable. It functions as a filter that modulates what receives further processing. As such, it is contrasted with Gregory’s (Gregory, 1974), Neisser’s (Neisser, 1976) and Churchland’s “strong interactionism,” which argues that top down processes affect the perceptual processing it self.

Research in linguistics shows weak interactionism to be quite plausible, especially if it is augmented by the assumption that the bottom up processes propose their alternatives in parallel (Forster, 1979; Crain & Steedman, 1985; Altmann & Steedman, 1988, Ferstl, 1994). The conceptual information that is brought to bear upon the syntactic or the phonological processor must be available after these processors have formed some hypotheses regarding the syntactic structure. Then conceptual information intervenes to select the appropriate structure in the specific context. The same model applied to vision may be able to account for most, if not all, of the cases that strong interactionism cite as evidence in its favor.

I will address now some of Churchland’s arguments regarding illusions. In the Müller-Lyer illusion the general assumptions regarding the projection of three-dimensional objects onto two dimensions deceive us into seeing the two lines as different in length. Fodor used this case to argue that, since we cannot help seeing the lines as unequal, despite our knowledge to the contrary, the visual input system is informationally encapsulated. Churchland (1988) answers that using this illusion to support impenetrability is a poor choice, because as Fodor admits, children with less experience with edges and corners are less susceptible to the illusion. But this means that this illusion is the result of learning from experience, which shows the cognitive penetrability of perception. There are three retorts against Churchland. First, it is by no means certain that this illusion is the result of misplaced depth constancy. Second, even if it is, and perceptual learning is involved, as it will be argued in the section regarding plasticity, perceptual learning does not necessarily involve cognitive top-down penetrability but only data-driven processes. Third, there are other, low-level explanation of the illusion, based on filtering interactive models.

The duck-rabbit and the Necker cube configuration exemplify illusions in which a figure or a scene is so ambiguous that its resolution seems to rely on higher-order cognitive factors. The context may dictate whether the rabbit-duck configuration is perceived as the picture of a duck or a rabbit, or the way the cube is perceived. Thus, perception is shown to be cognitively penetrable. This need not be the case, however. The weak interaction model can account for this phenomenon. The bottom up processes of vision propose in parallel both duck and rabbit, and higher cognitive states select one. Since the production of the output of the visual input systems is not affected by any top down processes, their impenetrability is not undermined.

Fodor (1988) offers another reply, which is worth discussing. He argues that one does not get the duck-rabbit configuration to flip by changing her assumptions and that “believing that it’s a duck doesn’t help you see it as one.” What does the trick is fixation at the appropriate parts of the configuration. Whether Fodor is right or wrong is a matter of empirical investigation. Some research (Taddei-Ferretti et al., 1996) renders justice to some of Fodor’s claims, but also suggests that top-down processes may determine the way one perceives reversible figures.

It has been found (Kawabata, 1986; Magnuson, 1970; Peterson & Gibson, 1991, 1993) that subliminal stimuli that reach perception but not cognition affect the reversal rate of reversible configurations, confirming the effect of bottom-up information on the perception of such ambiguous patterns. This research also shows that fixation at some crucial “focal areas” of the ambiguous pattern may cause reversion of the schema, thus justifying Fodor’s claim that fixation may do the trick. But research on reaction times in responses to such patterns shows that perception of the pattern is influenced by the previous interpretations of the incoming sensory information, suggesting that background information affects the way the ambiguous pattern is perceived. It may be, therefore, that specific-object information affects in a top-down manner the perception of some ambiguous patterns.

Churchland’s (Churchland, 1988) favorite example, finally, is the case of “inverted fields” in which people who wear lenses that invert the retinal image, after a period of disorientation, succeed in adjusting to the new perceptual regime and behave quite normally. According to Churchland, this is a clear demonstration of the penetrability and plasticity of our vision, since some very deep assumptions guiding the computations performed by the visual system can be reshaped in view of new experiences. He thinks that the adaptation to the inverting lenses shows the plasticity of some deep assumptions implicit in visual processing, such as the specific orientation of the visual field. These assumptions can be reshaped in a week or two, a fact that shows the plasticity and penetrability of perception.

Fodor (1988) concedes to Churchland that if in fact the case of inverting lenses reflects the penetration of perception by experience, then the modularity thesis is blown to pieces, since Fodor is committed to the thesis that the perceptual systems have a fixed neural architecture. Fodor, however, denies that this is the case, since one might expect to find this kind of plasticity on good ecological grounds. He claims, that there may be specific mechanisms that "function to effect the required visual-motor calibrations" that are required for the adaptation of the subjects to the new inverted retinal images. Fodor concedes to Churchland the plasticity of perception in this case, but he does not think it damaging for the thesis of the theoretical neutrality of observation, since it is expected on specific ecological grounds, whereas, the thesis of the theory-ladenness of observation requires that the perceptual field be reshaped "by learning physics."

There are several reasons why this may not be the case. First, the weak interaction model seems to work here too. In fact, the later stages of inverted field learning seem more like an oscillation between two competing impenetrable modules than anything penetrable. Second, Churchland thinks that the inverting lenses cause true visual effects. This is disputable, however, as there are accounts of this phenomenon that explain the reorientation and adaptation, not by imposing changes in the visual computational mechanisms and the visual representation itself, but by positing that such adaptations occur in the felt position of the limbs and the head (Dodwell, 1992).

Suppose, finally, that some of the basic assumptions underlying visual processing do change as a result of experience. Does this mean that perception is cognitively penetrable by top-down processes, and that people who wear inverting lenses readapt because they know that they see the world up-side down? Or do they readapt because, their perceptual mechanisms readjust? All the phenomenon shows is that perception is data-driven, not theory-driven, and that, in view of some new experience that upsets its balance, the system readjusts to achieve equilibrium. Though this undermines the claim of the rigidity of the modules' architecture, it does not imply that information flowing from higher cognitive levels penetrates the visual circuits to effect the appropriate changes. As we shall see later on that both Churchland and Fodor confuse cognitive penetrability with experience-driven changes in the perceptual systems. Evidence for the latter does not imply the former thesis.

2.1.4. Discussion

We have seen that nearly all cases of visual illusions reported by Churchland as evidence for the cognitive penetrability of perception can be explained away by other means other than cognitive penetrability. Weak interactionism figures prominently among these other alternative explanations.

There is, however, some evidence (Marslen-Wilson & Tyler, 1987; Marslen-Wilson, Tyler & Coster, 1993, for language Taddei-Ferretti et al., 1996, for vision) that the weak interactive model may be wrong, and that background information may play an active role in syntactic or visual disambiguation. Thus, one must consider the possibility that some top down flow of information may shape the output of the perceptual modules, undermining their cognitive encapsulation. What does this mean for perception? The answer to that can be found in Marr's view that it is only when the program of vision fails to resolve a particular ambiguity that a recourse to "object-specific" knowledge is needed.

The idea is that the top-down routes from higher cognition to vision may mediate processing within the modules, but are not on-line. The access to object-specific information occurs only on those occasions that a particular ambiguity cannot be resolved in a purely bottom-up manner. This conception of the cognition-perception links accounts for the almost cognitively-impenetrable character of vision, and is in concert with all the arguments that support such a view, while reconciling them with evidence for top-down flow of information.

Fodor could agree with such an approach. He has already argued along this line in discussing the modularity of the syntactic parser. Research (Marslen-Wilson & Tyler, 1987; Marslen-Wilson, Tyler & Coster, 1993) with open referents showed that the syntactic processor must have access to nonsyntactic information. Fodor, Garrett, Swinney (1993) reacted by modifying the theory of the informational encapsulation of the syntactic processor while retaining the basic idea of the independence of this processor. According to the modification, the syntactic processor is insulated from extra syntactic information until it meets an open referent. Then the processor opens to semantic information so that it can capture the correct filler. This syntactic-semantic interaction, however, does not take place on-line, but it must await the appearance of an open referent, which activates the top-down flow of information.

This Fodorian concession should not be confused with the weak interactive model, according to which the parser sends up to all possible syntactic interpretations and the context (specific world knowledge) determines which one will be chosen as fitting the context. This time Fodor allows for top down processes to affect the output of the syntactic module, since extralinguistic information intervenes to determine the appropriate filler so that the parser can continue its course. Even if this interaction is weak, that is, even if the parser sends up all possible fillers and the context selects the appropriate one, still the rest of the parsing process is cognitively contaminated. Thus the first pass processing is influenced by cognitive factors. Encapsulation however, is not threatened, since these communicative channels are off-line, that is, they are activated only when triggered by the presence of an open referent. If they were on-line, then inferential activity guided by the context would occur as the fragment sentence was being heard. According to Fodor, Garrett, and Swinney (1993) this is not the case (for a criticism of this thesis see Peatfield and Marslen-Wilson, 1995, who claim that semantic inferential activities do occur on-line while the sentence is being heard). In view of this, we can say that input modules are semiencapsulated from higher cognitive centers.

Perception, despite its bottom-up character, is not insulated from knowledge. There are three ways in which knowledge intrudes on perception. First, we have seen that even perceptual processes are informed and constrained by some general world principles that reduce indeterminancies in information. Second, in cases of ambiguities, perception is influenced by a top-down flow of information, which, however, intervenes only to select the hypothesis that fits the context, from among the hypotheses that perception produces insulated from any cognitive influences (weak interactionism). Since this is a postperceptual effect, it does not undermine the cognitive impenetrability of perception. Finally, there may be some rare cases in which information from higher cognitive centers intervenes and modulates the production of perception's output, not just the product of perception. These channels of top-down information, however, are not on-line.

Now, if Fodor believes that object recognition is impenetrable, then he is wrong and Churchland is right to point this out. Fodor (1983, 97) argues that the “visual-input system delivers basic categorizations.” He believes that the *3D* sketch is computed algorithmically in a bottom-up fashion and that an encapsulated processor can perform object identification. Then, at a final stage, the *3D* sketch is paired with a “form-concept’ dictionary, which selects one of the outputs of the encapsulated processor. Thus, Fodor seems to think that object identification involves weak interactionism, and that, in this sense, it is a purely bottom process.

Against that, it has been argued that the formation of the object-centered representation relies on top-down flow of information. – Pylyshyn (1999) argues, in accordance with Marr (1982), that the end product of perception is the *21/2D* sketch, which does not deliver objects but representations of shapes (structured representations of *21/2D* surfaces of objects). Thus, the output of the encapsulated visual processor does not consist in the regular objects that we encounter in the world, but in a more abstract form of categories of such objects, classified according to their generic shapes.

Fodor’s distinction between “fixation of appearances” or “observation,” which is the result of the functioning of the perceptual modules, and “fixation of belief” is misguided. He seems to distinguish between the “sensory” and “cognitive” or “semantic” processes that are involved in the formation of observation statements and considers observation as a precognitive activity whose output is processed by cognition giving rise to the observation statements. Philosophers would recognize here the distinction between what we see and how we interpret it. But this distinction is misleading, because object recognition is a cognitive process, and observation involves object-recognition. The distinction Fodor wishes to draw between a bottom-up, theory-neutral, and a top-down, theory-laden, process should not be cut at the observation/cognition interface, since such an interface does not exist, but at the perception/cognition interface. This criticism of Fodor’s views does not imply the rejection of the modularity thesis. What is rejected is the view that observation is informationally encapsulated.

Our discussion suggests that Churchland is right in rejecting the encapsulation of observation. He is wrong, however, in arguing for the cognitive penetrability of perception. First, it seems likely that, as the weak interaction model states, the higher cognitive states affect the products of the visual modules only after the visual modules have produced their product, by selecting, acting like filters, which output will be accepted for further processing. Weak interactionism can explain much, if not all, of Churchland’s evidence for strong interactions. Second, as we shall see later, perceptual learning need not involve top-down penetrability of perception from higher cognitive states. Finally, even if weak interactionism is wrong and there is some top-down flow of information from higher cognitive levels, this flow is not on-line and the input systems are generally cognitively impenetrable.

This ensures the neutrality of perception in real-world human activities. One would not argue that the rare cases (if any) of perceptual ambiguities that are resolved by means of cognitive interference imply that perception is cognitively penetrable, and that the scientists’ beliefs shape what they perceive. The point is that the results of scientific enterprise very rarely result in genuine perceptually ambiguous images, whose resolution requires top-down processes. Scientists may argue over what objects they see (object recognition), but they very rarely disagree as to what they perceive (the image extracted from early vision.

Moreover, one does not choose which arrays are ambiguous, nor does one choose the terms of ambiguity. Recall that these illusions arise exactly because they play against the physical assumptions underlying our vision. One does not choose these assumptions, they reflect our world. It follows that, although one may "choose" to see a rabbit, or a duck in a duck-rabbit configuration, one cannot "choose" one's illusions. This means that disagreeing scientists cannot argue that they perceive different things, due to the top-down character of perception, because the cases of real perceptual ambiguities that require top-down processing cannot be chosen as it fits their current dispute. This justifies Fodor's thesis that there is a theory-neutral basis, though it is not embedded in an observational, but rather in a perceptual vocabulary.

2.2. The argument from descending pathways

Churchland's (Churchland, 1988) second line of attack against the impenetrability of perception consists in evidence that there are top-down links from higher cognitive centers to the peripheral modules (here Churchland uses the term "module" in the way it is employed in neuroscience, according to which brains are structures, with cells, columns, layers and regions that divide up the labor of information-processing, and each such region is called a module). He reports findings from cell-staining techniques, according to which the ascending pathways from the retinal to the geniculate nucleus (LGN) and from there to the visual cortices and to other centers higher in the processing hierarchy are matched by descending pathways from the higher level of processing back to the earliest processing systems at the retina (Zeki, 1978; Van Essen, 1985). Churchland argues that the function of these descending pathways is "centrifugal control," that is, "the modulation of lower neural activity as a function of the demands sent down from levels higher in the cognitive hierarchy."

The existence of descending pathways is hardly disputable, nor is the importance of top-down processes in observation. What is disputable is the function of these top-down pathways in perception. Does their existence imply that even the early levels of vision are penetrable by this top-down flow of information? The answer to this question is no. Not because, as Fodor (1988, 194) claims, "if there is no cognitive penetration of perception, then at least 'descending pathways' aren't for that" –the argument is apparently a poor one– but because we know something about the function of the neuroanatomical structures in the brain.

Leopold and Logothetis (1996) studied in animals the activity of neurons in areas ranging from the primary visual cortex, where retinal signals first enter the brain, to the area called IT, which is the very end of one fork of visual processing. Their study showed that in the primary cortex only 18% of neurons changed their response according to the image perceived by the animal. In areas corresponding to the midway of visual processing about one half of the neurons changed their response. In the IT area almost all neurons did. These findings suggest that most of the neurons in early processing report information that can be extracted from the information recorded on the retina and are not influenced by the higher cognitive centers.

But this is not the end of the story. What is the function of the descending pathways and why are they so widespread if the top-down influence to early vision consists either in filtering the output of the modules, or at maximum, in some off-line top-down flow of information? The answer comes from cognitive neuroscience. Research (Posner & Petersen, 1990; Posner & Carr, 1992; Kosslyn et al., 1993; Posner & Raichle, 1994; Heinze et al., 1994; Ziegler et al., 1997) with positron emission topography (PET) and event-related potential (ERP) provides a spatio-temporal picture (literally) of the brain of subjects while they are performing (a) bottom up processes, such as passive visual tasks (viewing on a screen strings of consonants, words, and pseudowords), (b) processes that require some top-down influences, such as active attention-driven tasks (searching visual arrays for thickened letters), (c) processes that rely heavily on top-down semantic processing (generating a use in response to a visual word), and (d) processes that are purely top-down, such as imagery. This picture sheds light on the role of top-down pathways. At the same time it supports the principle of modular design and the independence of the lower levels of visual processing from top-down influences.

In studies of passive visual tasks, subjects were asked to fix their gazes on a point in the middle of a monitor, in which four kinds of complex stimuli were to appear: false fonts, letter strings, pseudowords and words. PET scans provided pictures of the activation of visual areas in the brain during these tasks. The analysis of these pictures relied on the assumption that the visual stimuli consisted of four codes. First, the 'words' presented were complex collections of visual features, second, these features were aligned to form the letters of the English alphabet, third, some of the 'words' had forms that satisfied the rules of English language, that is, they were English words, and fourth, some of these words had meanings.

The responses observed were responses to some, or all, of the four codes. All four groups produced bilateral responses in multiple areas of the visual system. The subtraction of the PET images when the brain processes the visual features of the array from the PET images in semantic processing shows that only words and pseudowords produced characteristic responses in the inner surface of the left cerebral hemisphere, an area which is related to semantic processes. This suggests the existence of two levels of analysis in the visual system. The brain initially analyzes the visual features of the stimuli regardless of relationships to letters and words. At a second level, the brain analyzes the visual word forms. I hasten to note here that the fact that the subtraction of the PET images reveal an intense activity in the left hemisphere when semantic processing is taking place does not mean that semantic processing is localized at that area only. The method of subtraction only highlights areas that are activated in the one task but not in the other. It does not reveal the entire area that participates in semantic processing. In fact, we know that areas in both hemispheres are related to semantic activity

More interesting were the PET images obtained in the active attention-driven visual tasks and in tasks of visual imagery. In the active attention-driven visual tasks, subjects were presented with succession of images on a screen and were asked to react whenever some attributes (color, motion, form) were different from one image to another (focal attention groups). The passive control group were instructed to watch the screen without reacting. The divided attention group, finally, were instructed to react to any changes whatsoever in the images.

The PET images of the passive group showed activations of areas traditionally associated with registration of color, motion and form in the extrastriate cortex. The subtraction of the divided attention PET images from the focal attention PET images allows the isolation of the areas that compute the specific features of the focal attention groups. The results were clear. Attention enhances blood flow at the same areas which are activated during the passive tasks. Thus, the same areas that process information in the passive tasks, process information in the active attention tasks, only this time their activation is amplified. The subtraction of the PET images in the passive acts from the PET images in the focal attention tasks allows us to track those areas outside the visual areas that are also activated only during the focal attention tasks and not during the passive tasks. Indeed, there were found areas in the basal ganglia and the anterior cingulate gyrus (an area at the underside of the frontal lobe). These areas seem to be the sources of the amplification observed when attention is involved and it is likely that they constitute the attentional networks activated in the focal group conditions.

Similar results were obtained with the visual imagery tasks. There is evidence (Farah, 1984) that visual imagery activates the same brain areas as visual perception. Behavioral studies suggest that the processing of imagery and of visual perception share some mental operations. Studies with patients (Kosslyn, 1988; Farah, 1991) showed that the mental operations that support visual imagery are highly localized (are carried out in precise locations) and distributed in many different areas in the brain (in both hemispheres). Also, many of the neural systems at which mental images are generated are the same as those that are activated during visual perception. Neuroimaging studies confirm these results. The subtraction of PET images during passive control tasks from the PET images in imagery tasks and from the PET images in visual perceptual tasks shows similar activations in imagery and perception, especially in posterior areas.

The PET studies were complemented by ERP studies of subjects who view words and consonant strings. Certain areas in the brain are activated about 100 ms after the word or the string is presented. Since these areas are activated irrespective of the stimuli, it can be surmised that they are activated by the features that words and consonant strings share, namely, visual physical features. Differences in the responses to words and consonant strings started about 150 ms after the stimuli appears. This means that the brain registers the word form 50 ms later than the visual features. What is important to note is that the ERP study shows that the distinction between words and consonant strings is not fed back by other higher processing areas but arises at the posterior sites of the cortex.

In other ERP studies, subjects were asked to search for a thickened letter in letter strings. This is clearly an attention-driven task, in which one would expect to find some top-down, task-driven, processes. Records of the electrical activity during the search show that this top-down activity involves the same processing areas that are involved in computing visual features. But the search for the thickened letter causes activity in these same areas only about 200 ms after the stimulus (recall that the activity recorded when these sites register the visual features takes place 100 ms after the stimulus). Thus, the computations involved in the top-down attention driven tasks take place in roughly the same brain areas–the same electrodes are activated, (predominately right posterior areas)– in which the bottom-up registration of visual features occurs, with a time delay of about 100 ms. Finally, similar studies of subjects performing semantic tasks, such as generation of the use of a noun, showed that word meaning is registered about 250 ms after presentation of the stimuli, and some of the areas activated are the same with those areas activated when visual physical features are processed.

Let me redraw the picture. 100 ms after the presentation of the stimulus (letter strings) an extensive part of our brain responds to the physical characteristics of the visual array. 150 ms after the stimulus these features fuse to a single form, and about 200 ms after it the voluntary task-driven search is registered in the same areas that process the visual features. Thus, the top-down effects of attention are delayed in time, involve the same anatomical areas as passive perception, except that attention amplifies their recordings. Finally, about 250 ms after the stimulus, some of the same areas participate in the semantic processing of the input.

What do these PET and ERP findings suggest for our discussion? Tasks that require more or less top-down flow of information activate broadly the same areas that are needed to compute the purely bottom-up tasks. The active attention studies showed that when top-down processes occur, the activation of these areas is enhanced and the source of this amplification is higher areas in the brain. In order for the factors that cause this amplification to be transmitted to the lower areas certain descending pathways are required. The same conclusion can be drawn from the visual imagery studies. Visual imagery demands that activation originate from higher cognitive centers and descend to the visual cortex in which imagined images are formed. The descending pathways accomplish exactly this.

It should be noted here that the PET evidence is compatible with the filtering model of weak interaction, as it pertains to the role of top-down influences in solving ambiguities of images. The descending pathways that select the images proposed in parallel by the perceptual modules are expected to be activated during this processing. But this same evidence goes beyond the mere disambiguation of images, to examine what happens in processes that are top-down in a stronger sense. Imagery, for instance, is a characteristic “strong” top-down process, in that higher cognitive centers determine the image, they do not select it (there is nothing to select from in the absence of perceptual stimulus). The point is that imagery and perception share roughly the same processing sites, and this is independent of the claims of the filtering model.

The same conclusion can be drawn with respect to attention-driven tasks. If the filtering model claims that in attention-driven tasks, perception sends up in parallel all possible interpretations and attention selects the one relevant, then the filtering interaction model may be wrong. Because, how could it explain the amplification of the activity in the processing areas, which seems to emanate from attention networks that are higher in the cognitive hierarchy? This amplification implies that the “attention centers” actively intervene in the processes that take place in the relevant sites and do not simply select the outcomes of insulated processes. This is, however, evidence for task dependent top-down processes, not for influences by specific object knowledge.

One could argue on behalf of weak interactionism, however, that in the focus attention tasks, attention modulates processing by singling out locations and features that, consequently, are preferably searched. This can be done by enhancing the output of the salient feature detectors, by lowering firing thresholds (Egeth et al., 1984; Kahneman & Treisman, 1992; McCleod et al., 1991). The allocation of attention occurs in a preperceptual stage, and thus, does not threaten the cognitive impenetrability of perception (for an extensive discussion of the role of attention see Pylyshyn, 1999). Since attention enhances the salient output, one could argue that this is the cause of the observed amplification of activity in the processing areas. Though this is a plausible explanation and one should wait for further empirical investigation before deciding on this matter, there is evidence (Moore & Egeth, 1998; Shih & Sperling, 1996) that attention affects the priority of items, that is, that attention is directed to those items earlier than in inattentive tasks, rather than enhancing sensory qualities. If this is the case, then focused attention cannot be invoked to explain the observed amplification.

Tasks that require more or less top-down flow of information activate broadly the same areas that are needed to compute the purely bottom-up tasks. How does the brain do this? The answer is the reentrant connections or mappings among neurons in the brain (Edelman, 1987, 1992). These reentry connections map, as it were, the activity of any system onto the others and reciprocally, by allowing the transmission of information in all directions.

Recall now Churchland's argument. We know that there are many neural connections devoted to bringing information back to the sensory systems from higher cognitive centers. This constitutes evidence for the mediation of the output of the perceptual modules by information from higher cognitive states. But, as we saw, the descending pathways most likely have another role to play. The sensory systems are fed back information from higher centers, and signals from higher areas reenter the brain areas that had processed before the signals that were transformed by the higher centers to the reentrant new signals. The same areas that process in a bottom-up way sensory specific information are also involved in higher-level thought (voluntarily attention-driven search, imagery), except that in the latter case they are reentered in a top-down manner, and their activation is amplified. Hence, the existence of descending pathways does not imply the cognitive penetrability of perception.

2.3. The argument from the plasticity of the brain and perceptual learning

2.3.1. The plasticity of the brain, innatism, and modularity

Churchland's (1988) third argument against impenetrability comes from neuroscientific studies suggesting the plasticity of the visual cortex in particular, and the cortex in general. Its line is that plasticity goes against Fodor's view that the perceptual modules are endogenously specified, hard-wired, and have a fixed neural architecture. Instead it shows that these systems are developed over time in a highly plastic system and that this development is context-environmentally driven (by the characteristics of the input to the cortex). It is true that the cortex can support a variety of representations early in development, and that there are no areas of the cortex prespecified for a certain function. When inputs are rewired so that they project to a different region of the cortex from what they usually do, then the new region develops some of the properties of the original target (Sur, Pallas & Roe, 1990). When a piece of the cortex is transplanted to a new location it develops projections that characterize the new location rather its origin (O'Leary, 1993). Finally, when the areas for the normal higher functions are bilaterally removed in early infancy, regions at considerable distance from the normal site can take over their function (Webster, Bachevalier & Ungerleider, 1995).

All this evidence supports the "protocortex" hypothesis against its competitor "protomap" hypothesis. The latter states (Rakic, 1988) that the differentiation of the cortex into areas with specific functions is determined by intrinsic molecular markers, or by a prespecification of the proliferative zone. This implies that the areal specification of the cortex is innately determined and hard-wired (specified by interactions at the molecular level and independent of thalamic inputs). The former states (Killackey, 1990; O'Leary, 1993) that the protocortex is initially undifferentiated and is divided up into specific areas largely as a result of the thalamic input, which means that the specialization of the cortex is not innate but determined through the interactions of an undifferentiated cortex with the environment.

This is hardly the image of hard-wired, innate, perceptual systems, since the differentiation of the cortex into specific areas, or maps, is not prespecified. This is not the end of the story, however. First, all the research showing the plasticity of the cortex involves experiments in vitro. There is some evidence that the thalamus may know about its preferred cortical targets in realistic in vivo situations (Niederer, Maimon & Finlay, 1995). This means that the sites to which thalamic input is projected may be prespecified, and that the areas in the cortex that will become the perceptual systems are innately determined. Second, even though there is evidence that the representational differentiation of the cortex is not innate, that is, that there are no innate representations in the form of stable connection weights among the neurons, there is strong evidence that there are some powerful innate architectural and chronotropic constraints that restrict the way the cortex develops while interacting with the environment (Thatcher, 1992; Shrager & Johnson, 1995; Elman et al., 1996; Johnson, 1997).

The architectural constraints consist in (Elman et al., 1996, 27–30) the types of neurons, their response characteristics, their firing threshold, type of transmitter, nature of pre- and postsynaptic changes, number of layers, packing density of cells, degree of interconnectivity, whether the connectivity is excitatory or inhibitory, and finally, the way the various modules are connected together to form the whole brain. The chronotropic constraints refer to the timing of events in the developmental process, which often is under genetic control. Shrager and Johnson (1995) propose a model which shows how a progressively grown wave of plasticity that traverses the cortex may govern the way the cortex is parcelled into different functional maps.

Evidence for this kind of architectural and chronotropic innatism abounds. Johnson and Morton (1991) show that while face-recognition becomes progressively species-specific due to interactions with the environment, infants at birth consistently orient toward species-like faces as opposed to other attractive stimuli. This shows that there is a subcortical head-start to the process of face recognition. Research in language (Bates et al., 1992, 1994) suggests that while the child does not begin with innate representations for language, it begins with strong computational constraints that normally lead to the familiar pattern of cortex parcellation, when the child is exposed to its STE. These constraints include computational asymmetries between the right and left hemisphere, and differences in speed and type of processing.

Elman et al., (1996), and Johnson (1997) argue in favor of the existence of some chronotropic and architectural constraints on cognition, while they reject the notion of some prespecified representations. They consider innate representations, however, as constant connection-synaptic weights. That does not preclude the possibility of representational innateness, in the form of predetermined initial connection weights, which will reconfigure, up to the limits of initial constraints, with learning.

We know that the development of axons and dendrites, as well as cell death, are postnatal phenomena, and that early in life there is an excess of synaptic connections among neurons. We also know that synaptic changes are related to the information received from the environment, in that those cells that participate in information processing are selected to survive, whereas the inactive ones become eliminated (Changeux & Dehaene, 1989; Kolb, 1993). However, there is evidence (Ebbesson, 1988) that the pattern and timing of neural connection loss is independent of experience and endogenously specified.

Johnson and Karmiloff-Smith (1992) proposed a model of synaptic loss and stabilization that purports to take into account evidence suggesting both the experience-dependent and experience-independent and time-constrained character of postnatal neural development. They distinguish between "timing" and "patterning" of loss and argue that although the specific pattern of synaptic loss is experience-dependent, the timing and extent of this loss is intrinsically determined. According to this view, the basic laminar structure of the cortex is consistent and is relatively insensitive to experience, whereas the differentiation in areas that characterizes the cortex is partly the product of experiential input.

In the same vein Johnson (1997, 40–41) argues that "while the basic architecture of a network is innate (basic circuitry, learning rules, type, and number of cells etc.), the detailed patterns of (dendritic and synaptic) connectivity are dependent upon experience. In such a case we may say that while the network imposes architectural constraints on the representations that emerge within it, there are no intrinsic representations." –Examples of the way some architectural biases can affect and restrict representations and learning can be found in O'Reilly and Johnson's (1994) model of chick-imprinting, and in Jacobs, Jordan and Barto's (1991) model of the "what" and "where" system. The idea is that the macrocircuitry of the brain is prespecified, whereas the detailed microcircuitry is the product of the interaction of the brain with the environment.

Churchland is right that the cortex can support a variety of representations early in development, and that there are not areas of the cortex prespecified for a certain function. This plasticity, furthermore, is not just a system response to early injury but is a basic property of neural operations (Antonini et al., 1995; Katz et al., 1995; Polat & Shagi, 1995). But evidence also suggests that the brain is not equipotential for all functions, and that there are deficits, due to brain damage, that persist to some extent. This shows that the brain is not entirely plastic.

If the early brain damage is such that it affects the spatial-analytic processing, a subtle deficit persists, unlike the cases concerning early brain injuries that affect language, in which the initial deficit is almost reversible (Stiles, 1995). This suggests that the neural system for processing spatial-analytic information is much more highly prespecified than the system processing language, and less plastic in its response to early injury. The reason may be that since spatio-analytic processing is phylogenetically much older than language processing, evolution has imposed a more rigid structure on the corresponding neural system. A parallel argument can be made for the visual system, and consequently, one could argue that visual perceptual processing is much more constrained, and much less plastic than the higher, phylogenetically younger, higher cognitive processes.

It seems, thus, that the brain is initially undifferentiated representationally but possesses strong architectural and chronotropic restrictions which restrict its representational potential and modulate the process of functional parcellation. This, conjoined with the view that there is a progressive loss of synapses due to task specialization, could explain the finding that in early phases of normal language development widely distributed regions of the brain participate in the processing (Mills, Coffey-Corina & Neville, 1993), whereas children who are efficient users of language show a much more focused distribution of brain electrical activity.

Where does the above discussion leaves us with respect to the problem of innateness of our perceptual systems? In view of the above, Karmilloff-Smith's (1992, 4) "minimalistic" stand toward innateness, which is also shared by Clark (1993a; 1993b) that " a fairly limited amount of innately specified predispositions (which are not strictly modular) would be sufficient to constrain the classes of inputs that the infant mind computes" seems quite plausible. These predispositions include some maturationally constrained attention biases, and some structures that constrain the way the input is processed. The former ensure that some part of the input (which is characteristic of a certain domain) will be processed by a mechanism that is suited to do exactly that and ignore other data, and the latter constrain the way these data are processed.

In order to see what the above means, let us turn to the function of the auditory system with respect to the acquisition of language. Since the infant's exposure to speech-input is part of her species typical environment (STE-Johnson, 1993) we can say that there are cells, with their dendrites and axons, which are sensitive to receiving the correct sensory inputs, and which, in response to the processing demands that this input imposes, reconfigure themselves to handle the task. These are the cells and connections that are selected to survive. This process of selection in response to input typical to the species is called experience-expectant (Greenough et al., 1993). Once this is done, these neural connections are sensitive only to a certain class of input (some particular speech sounds, as opposed to all kinds of sounds), process only this kind of input, and are indifferent to other sensory information.

The conclusion is that far from being an equipotential medium, the brain starts with some strong innate constraints that restrict the kinds of representations it builds, in response to external environmental input. We also see that the input modules (and the linguistic one) result from development and interaction with the environment rather than being there from the beginning. Thus, current research supports a view of progressive modularization with development that encompasses, in addition to perceptual modules, higher cognitive centers as well.

In view of this one should amend the claim made earlier that there are no domain-specific areas in the brain and allow the possibility that evolution has imposed on certain areas such strong constraints that, under normal conditions, these areas and these areas alone will perform certain specific functions. Architectural and chronotropic constraints seem to suffice to explain the parcellation of the brain. Jacobs, Jordan and Barto's (1991) "What and Where" connectionist system exemplifies this by showing how the architectural constraints of systems suffice to perform task decomposition and allocate one part of the system to one task, and the other to the other task, creating a complex modularized system consisting of two modules, each of which is domain-specific.

Churchland raised the issue of cortical plasticity to argue that there are no innately specified, domain-specific cortical areas, and thus, that the perceptual systems are neither hard-wired, nor encapsulated, since the idea of cortical plasticity does not "sit at all well with a picture of endogenously specified assumptions unique to each modality." (Churchland, 1989, 267). Against this claim, I discussed the kind of innatism and hard-wiring that is supported by current neuroscientific and neuropsychological research, and concluded that the brain develops and functions under certain architectural and chronotropic constraints that restrict the kind of representations it builds when interacting with the environment. The modules are the result of the interaction of a constrained brain with environmental input, and consequently, they appear with development and may become hard-wired, though initially they were not. But this is Fodor's (Fodor, 1983) view of innateness, and thus, Churchland's criticism fails.

2.3.2. Perceptual learning: evidence for cognitive penetrability?

Let us turn now to the issue of perceptual learning and consider whether it implies cognitive penetrability, as Churchland has repeatedly argued in his discussion of illusions. Is there a way to reconcile the notion of cognitive impenetrability of perceptual systems with the growing evidence for their diachronic penetrability? Fodor (1988) thought that there is not any, and that this issue would be resolved with the findings of empirical research. Should empirical research show perceptual learning to be possible, then the encapsulation of his input modules would have been proved false. Well, these findings are out and suggest that perceptual systems are diachronically, meaning in the long run, open to some rewiring of the patterns of their neural connectivity, as a result of learning. In other words, these systems are to some extent plastic. But this does not mean that they are cognitively penetrable. Let us see why.

Research by Karni and Sagi (1995), Yeo, Yonebayashi, and Allman (1995), Antonini, Strycker, and Chapman (1995), Stiles (1995), Merzenich and Jenkins (1995) show that changes can be induced in visual cortical neural patterns in response to learning. More specifically, visual processing at all levels may undergo long-term, experience-dependent changes. The most interesting form of learning is "slow learning," because it is the only type that causes structural changes in the cortex (formation of new patterns of connectivity). Such learning can result in significant performance improvement –for example, one may learn with practice to perform better at visual skills involving target and texture discrimination and target detection, and to learn to identify visual patterns in fragmented residues of whole patterns (priming). Performance in these tasks was thought to be determined by low-level, stimulus-dependent visual processing stages. The improvement in performance in these tasks, thus, suggests that practice may modify the adult visual system, even at the early levels of processing. As Karni and Sagi (1995, 95–6) remark "[L]earning (acquisition) and memory (retention) of visual skills would occur at the earliest level within the visual processing stream where the minimally sufficient neuronal computing capability is available for representing stimulus parameters that are relevant input for the performance of a specific task."

Karni and Sagi (1995) suggest that slow learning is independent of cortico-limbic processing, which is responsible for top-down processes and, through the interaction of the limbic system with the visual pathways, responsible for conscious object recognition (Webster, Bachevalier and Ungerleider, 1995). It is also independent of factors involving semantic associations. Yeo, Yonebayashi and Allman (1995) suggest that priming facilitates the neural mechanisms for processing images and that the cortex can learn to see ambiguous patterns by means of experience-induced changes in functional connectivity of the relevant processing areas. Thus, priming involves a structural modification of basic perceptual modules. Practice with fragmented patterns leads to the formation of the "priming memory" which may be stored in the cortical visual areas. Long-term potentiation (LTP) may be the mechanism implementing these architectural changes by establishing experience-dependent chains of associations.

Slow learning-induced architectural modifications are experience dependent, in that they are controlled by the "image" formed in the retina. But, although learning and its ensuing functional modifications occur in those neuronal assemblies that are activated by the retinal image, still some extraretinal factor should provide the mechanism that will gate functional plasticity. The problem is that although many neuronal assemblies are activated by the retinal image, learning occurs only in those assemblies that are behaviorally relevant. This is called the "gating of neuronal plasticity." The mechanism that ensures this relevance cannot be the retinal input.

The factor that modulates gating is the demands of the task. They determine which physical aspects of the retinal input are relevant, activating the appropriate neurons. Functional restructuring can occur only at these neuronal assemblies. The mechanism that accomplishes this is attention. Focusing attention ensures that the relevant aspects of the input be further processed. Attention intervenes before the perceptual processes (Pylyshyn, 1999) and determines the location at which search will be conducted and/or the relevant features that will be picked-up, since focal attention may enhance the output of the salient feature detectors by lowering firing thresholds (Egeth et al., 1984; Kahneman & Treisman, 1992; McCleod et al., 1991). There is indeed ample evidence for the necessary role of attention in perceptual learning (Ahissar & Hochstein, 1995) and for the role of attention in learning to perceive ambiguous figures (Kawabata, 1986; Peterson & Gibson, 1991, 1993)

Thus, slow learning takes place under specific retinal input and task-dependent conditions. These factors do not involve cognitive influences, and therefore, do not imply the cognitive penetrability of perception. Recall that slow-learning is independent of recognition and semantic memory. Most of the priming effects are associated with identification and discrimination of relative spatial relations and extraction of shapes. This brings to mind Hildreth and Ulmann's (1989) intermediate level of vision, mentioned in section 1. 1. The processes at this level (the extraction of shape and of spatial relations) are not purely bottom-up, but do not require the intervention of specific-object knowledge. They require the spatial analysis of shape and spatial relations among objects, which is task dependent.

So, the perceptual systems are to some extent plastic, as Churchland argues. But this plasticity is not the result of the penetration of the perceptual modules by higher cognitive states, but rather, the result of learning-induced changes, that are modulated by the retinal input and task-demands. Fodor (1988), given his view that the perceptual modules have a fixed architecture, had to concede that if evidence is found for diachronic changes in the functional architectures of the modules, then this would mean the end of the modularity thesis of perception. In the view of perceptual modules presented here, this is not necessarily so, since the data-driven changes in the perceptual modules can be accommodated by the notion that the modules are semihardwired. All this view requires is that the functional changes reshape the microcircuitry and not the macrocircuitry of the modules. Bearing in mind that priming enhances performance in various visual tasks, one cannot see how such learning could violate the general assumptions reflecting our world, and thus reshape the basic macrocircuitry.

3. Conclusion

My arguments suggest the picture of a modular organization of our perceptual systems, which is somewhat different from that of Fodor. According to this picture, which is based on the distinction between noncognitive perception and cognitive observation, the perceptual semi-Fodorian modules share some properties with the Fodorian ones. They are domain-specific, fast, automatic, mandatory, independent of conscious control, and emerge as the result of the interaction between a constrained brain and the environment. They are different from their cousins in that, first, they are informationally semiencapsulated. The term “semi” is used to convey my notion of off-line penetrability, should the information available in the visual array in combination with the general theory upon which visual computations rely prove insufficient to resolve ambiguities.

Second, the modules have no fixed neural architecture, but are semi hard-wired, since perceptual learning can modify the patterns of connectivity. Fodor’s insistence that the modules have a fixed architecture made him vulnerable to Churchland’s attacks. It also committed him to allowing that if perceptual learning occurs, then modularity fails. As we saw, perceptual learning is possible and restructures the visual system. But, this does not affect his claim with regard to the module’s cognitive on-line impenetrability.

We see the problem in Churchland’s arguments. First, his argument regarding reentrant pathways does not imply by itself a top-down flow of cognitive, that is, extravisual, information. It seems most likely that the top-down pathways serve to convey information to locations along the route of early vision, or perception, so that these locations can participate in processing of cognitive tasks and not in order to carry out their perceptual function.

Attempting, second, to demonstrate that the thesis about the cognitive penetrability of our perception, and the theoretical neutrality of observation is false, he confuses the plasticity of the brain and perceptual learning with cognitive penetrability. But the former does not entail the latter. The only way out for Churchland is to argue that the experience-induced learning changes the way we observe the world, and this, in its turn, by means of some top-down flow of information which affects the way we perceive. Though it is true that our experiences shape our theories and the way we see the world, to say that these theories influence the way we perceive the world is question begging, since one must show that this top-down influences occur. Fodor seems to allude to this when he argues that even though the perceptual modules may be plastic and amenable to local functional changes due to environmental pressure, still “these might not even be perceptual effects of acquiring beliefs; perhaps they’re perceptual effects of having the experiences in virtue of which the beliefs are acquired” (Fodor 1988, 192).

I have drawn a picture of perceptual modules that are semi-Fodorian. They are semipenetrable (there may be some off-line top-down influences), innate, semihard-wired, domain-specific, fast, automatic, and mandatory. Where does this leave us regarding the theoretical neutrality of perception? Fodor bases his claim on the modularity of perception. Churchland reacted by arguing that perception is cognitively penetrable, since it is influenced in a top-down manner by higher cognitive states, and thus it cannot be theory-neutral. Bearing in mind the distinction between perception and observation, one can argue that although Churchland is right that observation is cognitively penetrable, perception may be cognitively semi-impenetrable and, thus, theory neutral.

This theoretical neutrality does not mean that the perceptual modules function independent of theories. Perceptual computations are based on some general assumptions about the world that constitute a powerful theory constraining visual processing. But this theory is not learned. It is a prerequisite, if learning is to start off the ground, and is shared by all. Thus, this theoretical ladenness cannot be used as an argument against the existence of a common ground (in the form of some innate principles predetermined by the structure implicit in the interaction of the perceptual mechanisms with the outside world), because perception based on a theory that is shared by all, is such a common ground.

Notes

1. All references to Fodors article are from the reprint in A. I. Goldman (ed.) *Readings in Philosophy and Cognitive Science.*
2. All references to this article are from the reprint, as Chapter 12, in Churchland (1989, 255–279).
3. It should be emphasized that these object recognition units are not necessarily semantic, since we may recognize an object that we had seen before, even though we have no idea of its name, of what it does and how it functions, that is, even if we have no semantic and lexical information about it. Thus, Humphreys and Bruce (1989) introduced a distinction between the perceptual classification and semantic classification and naming. These processes are independent one of the other (as we will see when we discuss the various forms of visual agnosias).

References

Ahhisar, M., & Horchstein, S. (1993). Attentional control of early perceptual learning. *Proceedings of the National Academy of Science,* USA, *90,* 5718–5722.

Altmann, G. T., & Steedman, M. (1988). Interaction with context during human sentence processing. *Cognition, 30,* 191–238.

Antonini, A., Strycker, M. P., & Chapman, B. (1995). Development and plasticity of cortical columns and their thalamic inputs. In B. Julesz and I. Kovacs (Eds.), *Maturational windows and adult cortical plasticity.* Reading, MA: Addison-Wesley.

Bates, E., Thal, D., & Janowsky, J. (1992). Early language development and its neural correlates. In I. Rapin and S. Segalowitz (Eds.) *Handbook of Neuropsychology, 7,* Child Neuropsychology. Amsterdam: Elsevier.

Bates, E., Thal, D., Aram, D., Eisele, J., Nass, R., & Trauner, D. (1994). From first words to grammar in children with focal brain injury. Forthcoming in D. Thal and J. Reilly (eds.), *Special Issue on Origins of Communication Disorders, Developmental Neuropsychology.*

Biederman, I. (1987). Recognition by components: A theory of human image understanding. *Psychological Review, 94,* 115–147.

Bruce, V., & Green, P. R. (1993). *Visual perception: Physiology, pychology and ecology.* 2nd edition. Hillsdale, NJ: Lawrence Erlbaum.

Campion, J., & Latto, R. (1985). Apperceptive agnosia due to carbon poisoning. An interpretation based on critical band masking from disseminated lesions. *Behavioural Brain Research, 15,* 227–240.

Cavanagh, P. (1988). Pathways in early vision. In Z. W. Pylyshyn (Ed.), *Computational processes in human vision: An interdisciplinary perspective.* Ablex Publishing.

Changeux, J-P., & Dehaene, S. (1989). Neuronal models of cognitive functions. *Cognition, 33,* 63–109.

Churchland, P. M. (1988). Perceptual plasticity and theoretical neutrality: A reply to Jerry Fodor. *Philosophy of Science, 55,* 167–187.

Clark, A. (1993a). *Associative engines: connectionism, concepts, and Representational change.* Cambridge, MA: The MIT Press.

Clark, A. (1993b). Minimal nativism. *Mind, 102,* 587–610.

Crain, S., & Steedman, M. (1985). On not being led up the garden path: The use of context by the psychological syntax parser. In D. R. Dowty, L. Karttunen, & A. M. Zwicky (Eds.), *Natural language parsing: psychological, computational, and theoretical perspectives.* Cambridge: Cambridge University Press.

De Renzi, E., Scotti, G., & Spinnler, H. (1969). Perceptual and associative disorders of visual recognition. *Neurology, 19,* 634–642.

Dodwell, P. C. (1992). Perspectives and transformations. *Canadian Journal of Psychology, 46,* 510–538.

Ebbesson, S. O. E. (1988). Ontogenetic parcellation: Dual processes. *Behavioral and Brain Sciences, 11,* 548–549.

Edelman, G. M. (1987). *Neural Darwinism.* New York: Basic Books.

Edelman, G. M. (1992). *Bright air, brilliant fire. On the matter of mind.* New York: Basic Books.

Egeth, H. E., Virzi, R. A., & Garbart, H. (1984). Searching for conjuctively defined targets. *Journal of Experimental Psychology, 10,* 32–39.

Elman, J. L., Bates, E. A., Johnson, M. H., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1996). *Rethinking innateness: A connectionist perspective on development.* Cambridge, MA: The MIT Press.

Farah, M. J. (1984). The neurological basis of mental imagery: A component analysis. In S. Pinker (Ed.), *Visual Cognition* (pp. 245–271). Cambridge, MA: The MIT Press.

Farah, M. J. (1991). *Visual agnosia.* Cambridge, MA: The MIT Press

Felleman, D. J., Xiao, Y., & McClendon, E. (1997). Modular organization of occipital-temporal pathways: Cortical connections between visual area 4 and visual area 2 and posterior inferotemporal ventral area in macaque monkeys. *Journal of Neuroscience, 17,* 3185–3200.

Ferstl, E. C. (1994). The construction-integration model: a framework for studying context effects in sentence processing. *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society* (pp. 289–293). Hillsdale, NJ: Lawrence Erlbaum.

Ferster, D. (1981). A comparison of binocular depth mechanisms in areas 17 and 18 of the cat visual cortex. *The Journal of Physiology, 311,* 623–655.

Fodor, J. (1983). *The Modularity of mind.* Cambridge, Mass: The MIT Press.

Fodor, J. (1984). Observation reconsidered. In *Philosophy of Science, 51,* 23–43. Reprinted in A. I. Goldman (Ed.), *Readings in philosophy and cognitive science.* Cambridge, Mass: The MIT Press.

Fodor, J. (1988). A Reply to Churchland's Perceptual plasticity and theoretical neutrality. *Philosophy of Science, 55,* 188–198.

Fodor, J. A., Garrett, M., & Swinney, D. (1993). A modularity effect of parsing. Manuscript, Program in Philosophy, CUNY Graduate Center.

Forster, K. I. (1979). Levels of processing and the structure of the language processor. In W. Cooper and E. Walker (Eds.), *Sentence processing: Psycholinguistic studies presented to Merrill Garrett.* Hillsdale, NJ: Lawrence Erlbaum.

Gilchrist, A. (1977). Perceived lightness depends on perceived spatial arrangement. *Science, 195,* 185–187.

Goodale, M. A. (1993). Visual pathways supporting perception and action in the primate cerebral cortex. *Current Opinion in Neurobiology, 3,* 578–585.

Goodale, M. A. (1995). The cortical organization of visual perception and visuomotor control. In S. M. Kosslyn and D. N. Osherson (Eds.), *Visual cog*nition (Vol. 2). Cambridge, MA: The MIT Press.

Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in Neuroscience, 15,* 20–25.

Gregory, R. (1974). *Concepts and mechanisms of perception.* New York: Charles Scribners and Sons.

Greenough, W. T., Black, J. E., & Wallace, C. S. (1993). Experience and brain development. In M. H. Johnson (Ed.), *Brain development and cognition: a reader.* Cambridge: Basil Blackwell.

Heinze, H. J., Mangun, G. R., Burchert, W., Hinrichs, H., Scholtz, M., Muntel, T. F., Gosel, A., Schreg, S., Johannes, S., Hundeshagen, H., Gassaniga, M. S., & Hillyard, S. A. (1994). Combined spatial and temporal imaging of brain activity during visual selective attention in humans. *Nature, 372,* 543–546.

Hildreth, E. C., & Ulmann, S. (1989). The Computational study of vision. In M. I. Posner (Ed.), *Foundations of cognitive science.* Cambridge, MA: The MIT Press.

Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *The Journal of Physiology, 195,* 215–43.

Humphreys, G. W., & Riddoch, M. J. (1984). Routes to object constancy: Implications from neurological impairments of object constancy. *Quarterly Journal of Experimental Psychology, 36A,* 385–415.

Humphreys, G. W., & Riddoch, M. J. (1985). Authors' correction to "Routes to object constancy." *Quarterly Journal of Experimental Psychology, 37A,* 493–495.

Humphreys, G. W., & Bruce, V. (1989). *Visual cognition: computational, experimental and Nnuropsychological perspectives.* Hove, UK: Lawrence Erlbaum.

Jacobs, R. A., Jordan, M. I., & Barto, A. G. (1991). Task decomposition through competition in a modular connectionist architecture: The what and where vision tasks. *Cognitive Science, 15,* 219–250.

Johnson, M. H. (1997). *Developmental cognitive neuroscience.* Oxford: Blackwell.

Johnson, M. H. (1993). Constrains in Cortical Plasticity. In M. H. Johnson (Ed.), *Brain development and cognition: a reader.* Cambridge: Basil Blackwell.

Johnson, M. H., & Morton, J. (1991). *Biology and development: The case of face recognition.* Oxford: Blackwell.

Johnson, M. H., & Karmiloff-Smith, A. (1992). Can neural selectionism be applied to cognitive development and its disorders? *New Ideas in Psychology, 10,* 35–46.

Johnston, A. (1992). Object constancy in face processing: Intermediate representations and object forms. *Irish Journal of Psychology, 13,* 425–438.

Kahneman, D., & Treisman, A. (1992). The reviewing of object files: Object-specific integration of information." *Cognitive Psychology, 24:*2, 175–219.

Kawabata, N. (1986). Attention and depth perception. *Perception, 15,* 563–572.

Karmiloff-Smith, A. (1992). *Beyond modularity: A developmental perspec*tive *on cognitive Sience.* Cambridge, MA: The MIT Press.

Karni, A., & Sagi, D. (1995). A memory system in the adult visual cortex. In B. Julesz and I. Kovacs (Eds.), *Maturational windows and adult cortical plasticity.* Reading, MA: Addison-Wesley.

Katz, E., Victor, J. D., & Purpura, K. (1995). Dynamic changes in cortical response following removal and restoration of nearby visual inputs. In B. Julesz and I. Kovacs (Eds.), *Maturational windows and adult cortical plasticity.* Reading, MA: Addison-Wesley.

Killackey, H. P. (1990). Neocortical expansion: An attempt toward relating phylogeny and ontogeny. *Journal of Cognitive Neuroscience, 2,* 1–17.

Koch, C., & Poggio, T. (1987). Biophysics of computational systems: Neurons synapses, and membranes. In G. M. Edelman, W. E. Gall, and W, M. Cowan (Eds.), *Synaptic function*. New York: John Wiley and Sons.

Kolb, B. (1993). Brain development, plasticity, and behavior. In M. H. Johnson (Ed.), *Brain development and cognition: a reader.* Cambridge: Basil Blackwell.

Kosslyn, S. M. (1988). Aspects of cognitive neuroscience of mental imagery. *Science, 240,* 1621–1626.

Kosslyn, S. M., Alpert, N. M., Thompson, W. L., Malikovic, V., Weise, S. B., Chabris, C. F., Hamilton, S. E., Rauch, S. L., & Buonnano, F. S. (1993). Visual mental imagery activates topographically organized visual cortex. *Journal of Cognitive Neuroscience, 5,* 263–287.

Lawson, R., Humphreys, G. W., & Watson, D. G. (1994). Object recognition under sequential viewing conditions: Evidence for viewpoint-specific recognition procedures. *Perception, 23,* 595–614.

Leopold, D. A., & Logothetis, N. K. (1996). Activity changes in early visual cortex reflect monkeys' percepts during binocular rivalry. *Nature, 379,* 549.

Livingstone, M. S., & Hubel, D. H. (1987). Psychophysical evidence for separate channels for the perception of form, color, movement, and depth. *The Journal of Neuroscience, 7,* 3416–3468.

Magnuson, S. (1970). Reversibility of perspective in normal and stabilized viewing. *Scandinavian Journal of Psychology, 11,* 153–156.

Marr, D. (1982). *Vision: A computational investigation into human representation and processing of visual information.* San Francisco, CA: Freeman.

Marslen-Wilson, W., Tyler, L. K. (1987). Against modularity. In J. L. Garfield (Ed.), *Modularity in knowledge representation and natural understanding.* Cambridge, Mass.: MIT Press.

Marslen-Wilson, W., Tyler, L. K., & Koster, C. (1993). Integrative processes in utterance resolution. *Journal of Memory and Language, 32,* 647–666.

McLeod, P., Driver, J., Dienes, Z., & Crisp, J. (1991). Filtering by movement in visual search. *Journal of Experimental Psychology: Human Perception and Performance, 17,* 55–64.

Merzenich, M. M., & Jenkins, W. M. (1995). Cortical plasticity, learning, and learning dysfunction. In B. Julesz and I. Kovacs (Eds.), *Maturational windows and adult cortical plasticity.* Reading, MA: Addison-Wesley.

Mills, D. L. S., Coffey-Corina, S. A., & Neville, H. J. (1993). Language acquisition and cerebral specialization in 20-month-old infants. *Journal of Cognitive Neuroscience, 5:* 3, 317–334.

Moore, C. M., & Egeth, H. (1998). Perception without attention: Evidence of grouping under conditions of inattention. *Journal of Experimental Psychology: Human Perception and Performance, 23,* 339–352.

Neisser, U. (1976). *Cognition and Reality.* San Francisco, CA: Freeman.

Newcombe, F., & Ratcliff, G. (1977). Agnosia: A disorder of object recognition. In F. Michel and B. Scot (Eds.), *Les syndromes de disconnection calleuse chez l' homme.* Lyon: Colloque International de Lyon.

Niederer, J., Maimon, G., & Finlay, B. (1995). Failure to reroute or compress thalamocortical projections after prenatal posterior cortical lesions. *Society of Neuroscience Abstracts,* 21.

O'Leary, D. D. M. (1993). Do cortical areas emerge from a protocortex? In M. Johnson (Ed.), *Brain development and cognition: A reader.* Oxford: Blackwell.

O'Reilly, R. O., & Johnson, M. H. (1994). Object recognition and sensitive periods: A computational analysis of visual Imprinting. *Neural Computation, 6,* 357–389.

Rakic, P. (1988). Specification of cerebral cortical areas. *Science, 241,* 170–176.

Peatfield, M. D., & Marslen-Wilson, W. (1995). Pragmatic effects in zero anaphor resolution: Implications for modularity. In J. D. Moore and J. F. Lehman (Eds.), *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society* (pp. 364–369). Mahwah, NJ: Lawrence Erlbaum.

Perrett, D. I., Oram, M. W., Hietanen, J. K., & Benson, P. J. (1994). Issues of representation in object vision. In M. J. Farah and G. Ratcliff (Eds.), *The neuropsychology of higher vision: Collated tutorial essays.* Hillsdale, NJ: Lawrence Erlbaum.

Peterson, M. A., & Gibson, B. S. (1991). Directing spatial attention within an object: Altering the functional equivalence of shape descriptions. *Journal of Experimental Psychology: Human Perception and Performance, 17,* 170–182.

Peterson, M. A., & Gibson, B. S. (1993). Shape recognition contributions to figure ground organization in three-dimensional displays. *Cognitive Psychology, 25,* 383–429.

Poggio, G. F., & Talbot, W. H. (1981). Mechanisms of static and dynamic stereopsis in foveal cortex of the rhesus monkey. *The Journal of Physiology, 315,* 469–492.

Polat, U., & Shagi, D. (1995). Plasticity and interactions in early vision. In B. Julesz and I. Kovacs (Eds.), *Maturational windows and adult cortical plasticity.* Reading, MA: Addison-Wesley.

Posner, M. I., & Petersen, S. E. (1990). The attention system of the human brain. *Annual Review of Neuroscience, 13,* 25–42.

Posner, M. I., & Carr, T. H. (1992). Lexical access and the brain: Anatomical constraints on cognitive models of word recognition. *American Journal of Psychology, 105,* 1–26.

Posner, M. I., & Raichle. M. E. (1994). *Images of mind.* New York: The American Scientific Library.

Pylyshyn, Z. (1999). Is Vision continuous with cognition? *Behavioral and Brain Sciences, 22,* 341–365.

Shih, S. I., & Sperling, G. (1996). Is there feature-based attentional selection in visual search? *Journal of Experimental Psychology: Human Perception and Performance, 22,* 758–779.

Shrager, J. F. (1990). Commonsense perception and the psychology of theory formation. In J. F. Shrager and P. Langley (Eds.), *Computational models of sientific discovery and theory formation.* San Mateo, CA: Morgan Kaufmann.

Shrager, J. F., & Johnson, M. H. (1995). Waves of growth in the development of cortical function: A computational model. In B. Julesz and I. Kovacs (Eds.), *Maturational windows and adult cortical plasticity.* Reading, MA: Addison-Wesley.

Spelke, E. S. (1988). Object perception. In A. I. Goldman (Ed.), *Readings in philosophy and cognitive science.* Cambridge, MA: The MIT Press.

Stiles, J. (1995). Plasticity and development. Evidence from children with early occurring focal brain injury. In B. Julesz and I. Kovacs (Eds.), *Maturational windows and adult cortical plasticity.* Reading, MA: Addison-Wesley

Sur, M., Pallas, S. L., & Roe, A. W. (1990). Cross-modal plasticity in cortical development: differentiation and specification of sensory neocortex. *Trends in Neuroscience, 13,* 227–233.

Taddei-Ferretti, C., Musio, C., Santillo, S., Colucci, R. F., & Cotugno, A. (1996). The top-down contributions of will to the modulation of bottom-up inputs in the reversible perception phenomenon. In R. Moreno-Diaz and J. Mira-Mira (Eds.), *Brain pocesses, Theories and models: An international conference in honor of W. S. McCulloch 25 Years after his Death.* Cambridge, MA: The MIT Press.

Taylor, A. M., & Warrington, E. K. (1971). Visual agnosia: A single case report. *Cortex, 7,* 152–161.

Thatcher, R. W. (1992). Cyclic cortical reorganization during early childhood. *Brain and Cognition, 20,* 24–50.

Ulmann, S. (1979). *The Interpretation of visual motion.* Cambridge, MA: The MIT Press.

Van Essen, D. C. (1985). Functional organization of primate visual cortex. In A. Peters and E. G. Jones (Eds.), *Cerebral cortex* (Vol. 3): *Visual cortex.* New York: Plenum.

Yeo. R. M., Yonebayashi, Y., & Allman, J. M. (1995). Perceptual memory of cognitively defined contours: A rapid, robust and long-lasting form of memory. In B. Julesz and I. Kovacs (Eds.), *Maturational windows and adult cortical pasticity.* Reading, MA: Addison-Wesley.

Warrington, E. K. (1975). The selective impairment of semantic memory. *Quarterly Journal of Experimental Psychology, 27,* 635–657.

Warrington, E. K. (1982). Neuropsychological studies in object recognition. *Philosophical Transactions of the Royal Society, B298,* 15–33.

Warrington, E. K., & Taylor, A. M. (1973). The contribution of the right parietal lobe to object recognition. *Cortex, 9,* 152–164.

Warrington, E. K., & Taylor, A. M. (1978). Two categorical stages of object recognition. *Perception, 7,* 695–705.

Watt, R. J., & Morgan, M. J. (1984). Spatial filters and the localization of luminance changes in human vision. *Vision Research, 24:10,* 1387–1397.

Webster, M. J., Bachevalier, J., & Ungerleider, L. G. (1995). Development and plasticity of visual memory circuits. In B. Julesz and I. Kovacs (Eds.), *Maturational windows and adult Ccrtical Ppasticity*. Reading, MA: Addison-Wesley.

Zeki, S. M. (1978). Uniformity and diversity of structure and function in rhesus monkey prestriate visual cortex. *Journal of Physiology, 277,* 273–290.

Ziegler, J. C., Besson, M., Jacobs, A. M., & Nazir, T. A. (1997). Word, pseudoword, and nonword processing: A multitask comparison using event-related brain potentials. *Journal of Cognitive Neuroscience, 9:6,* 758–775.