

# The Phenomenal Content of Experience

ATHANASSIOS RAFTOPOULOS AND VINCENT C. MÜLLER

---

**Abstract:** We discuss in some length evidence from the cognitive science suggesting that the representations of objects based on spatiotemporal information and featural information retrieved bottom-up from a visual scene precede representations of objects that include conceptual information. We argue that a distinction can be drawn between representations with conceptual and nonconceptual content. The distinction is based on perceptual mechanisms that retrieve information in conceptually unmediated ways. The representational contents of the states induced by these mechanisms that are available to a type of awareness called phenomenal awareness constitute the phenomenal content of experience. The phenomenal content of perception contains the existence of objects as separate things that persist in time and time, spatiotemporal information, and information regarding relative spatial relations, motion, surface properties, shape, size, orientation, color, and their functional properties.

## 1. Introduction

Philosophers (Crane, 1992; Evans, 1982; Heck, Jr., 2000; Kelly, 2000; Peacocke, 1992, 2001; Smith, 2002; Shoemaker, 2002, Tye, 2002; 2005) have recognized the necessity of a kind of experiential content that is not conceptual and yet it is evaluable as correct or as incorrect. Having correctness conditions, it presents (Heck, 2000), or represents (Peacocke, 2001) the world as being in a certain way, though the representation does not have the structure of discursive judgments; it is not the content of beliefs.

Several reasons led to the endorsement of a nonconceptual content. First, the demand to account for how concepts/symbols are grounded in the world (Evans, 1982; Lowe, 1992; Peacocke, 1992; Shoemaker, 2002; Tye, 2002), ‘the symbol-grounding problem’. According to the descriptive theories of reference, a symbol is associated with a concept in the mind that constitutes its meaning and determines its reference, by allowing one to pick out the objects that are ‘described’ by the concept. Descriptive theories of reference by explaining references by descriptive means appeal to the descriptions of other symbols; that is, by appealing to concepts and the reference of other symbols. To escape the infinite regress, there would have to be some symbols whose reference does not depend on that of others and are founded directly in the world Devitt (1996, p. 159). Since perception provides us with an immediate contact with the world, the grounding should be sought in perception.

---

**Address for correspondence:** Department of Psychology, University of Cyprus, P.O. Box 20537, Nicosia, 1678 Cyprus.

**Email:** raftop@ucy.ac.cy

*Mind & Language*, Vol. 21 No. 2 April 2006, pp. 187–219.

© 2006 The Authors

Journal compilation © 2006 Blackwell Publishing Ltd

Second, nonconceptual content allows reference to objects while their features undergo changes—the problem of reference stability (Devitt, 1996; Pylyshyn, 2001). Suppose that the concept of an object is defined by a set of descriptions, and that some of the features of the object change. How could one assign the new properties to the same object? Relying exclusively on descriptions renders this impossible, since a change in the defying description implies a change in reference. If one could pick out objects by nonconceptual means, the problem would be solved. Since one's immediate access to one's environment is through one's experience, finding such a nonconceptual way of reference requires that at least part of the content of experience be nonconceptual. Pylyshyn (2001, p. 138) claims that this requires a 'demonstrative' reference, which picks out the object encoding none of its properties.

A third reason for the introduction of nonconceptual content is the necessity to account for the fine-grained character and richness of the 'phenomenal content' of experience (Crane, 1992; Evans, 1982; Heck, 2001; Peacocke, 1992; 2001; Tye, 2002; Shoemaker, 2002). The content of our experiences is usually much richer than any attempt to report it could be; it cannot be described conceptually. The phenomenal content of an experience is what is responsible for the 'phenomenal character' of experience. 'What it is like to have a visual experience (the "phenomenal character" of experience) is a matter of a certain sort of representational content that the experience has' (Tye 2002, p. 448).

There is a convergence of views as to the kinds of properties that figure in the list of the phenomenal properties of visual experience. Viewer-centered shapes, relative spatial relations, size, orientation, and location of objects, and colors and their properties (shades) figure in this list. Peacocke (1992) and Bermúdez (1995) include the representations of viewpoint-independent objects (Marr's 3D objects) in the list of nonconceptual properties.

There are several ways to put an item on the list of nonconceptual properties. One is by considering whether it is conceivable for a perceptual state to have a certain content without the perceiver exercising or possessing the concept articulating the content. This way is 'phenomenological', in that it relies on the introspection of experience in an attempt to pinpoint its nonconceptual part. One can see a certain shape without possessing the concept that conceptualizes that shape. However, one cannot see it as that (say, square) shape.

A second method is by considering cases in which some intentional behavior could be described as being caused by certain concepts (as when adults see an object disappearing behind a screen and because they know that objects do not simply vanish they decide to search for it) but we have reasons to believe that the organism that engages in such a behavior does not possess the requisite concepts or perhaps any concepts at all, if the *Autonomy Thesis* (Peacocke, 1992) holds (the thesis states that a creature can be in states with nonconceptual content, even though the creature possesses no concepts whatsoever). This is the case with infants or animals, which exhibit intentional behavior that in adults could be explained by reference to concepts, which the infants and the animals clearly do not possess

(children look for hidden objects without possessing the concept 'object'). The content of the states that the animals and infants have is nonconceptual.

A third way of determining nonconceptual content is by considering cases in which the content cannot be described conceptually because language does not have the requisite resources. Consider the large number of different shades that one can perceive. Few among them can be described conceptually; natural languages simply do not have the required conceptual arsenal. The content of many computational sub-personal states cannot be described conceptually. Natural languages do not conceptualize the contents of some of the early stages of visual processes. Thus, their content, whatever it might be, is nonconceptual.

These methods associate directly or indirectly perception with awareness and seek the phenomenal content of perception through the introspection of the content of perceptual states. We call them 'phenomenological methods'. They are beset by many problems. First, the awareness on which introspection relies has content that can be reported. Thus, the experiential content to which one has this kind of awareness is conceptual content. To unearth the nonconceptual content one has to think, and argue for, whether it would be possible to be in states with such contents if one did not possess or exercise the concepts that conceptualize those contents. But, as we shall see, conceptual content and nonconceptual content are different kinds of content.

Second, there is nonconceptual content of which one cannot be aware. Thus, the phenomenological method restricts the nonconceptual content to items of which one can be aware. As an example, consider the exclusion from the list of nonconceptual content of, say, the contours and boundaries of objects, which are nonconceptual and yet one cannot be aware of them, except as properties of objects at which one arrives by means of an abstraction that is a conceptual act. To alleviate this problem Bermúdez (1995) introduces a distinction between the nonconceptual content of perception and the nonconceptual content of subpersonal computational states. However, we shall see that his account has certain shortcomings.

A third problem of the phenomenological methods is the inclusion in the list of nonconceptual content of properties that are extracted from the scene by visual processes that are cognitively modulated and thus can only be conceptual contents. One imagines that one could perceive the 3D shape of an object even if one did not possess the concepts that correspond to that shape, and because of that one is led to believe that the 3D shape of an object may be the nonconceptual content of one's perceptual states. However, we shall see that the 3D shape of an object is extracted from the scene by cognitively modulated process and thus the 3D shape cannot be nonconceptual content.

In this paper we take the existence of a nonconceptual representational content of experience for granted and offer an alternative approach for determining it. We propose to specify the phenomenal nonconceptual content via a study of the perceptual mechanisms that deliver that content. The reason is that if there are no mechanisms that retrieve information from a visual scene in bottom-up

conceptually unmediated ways—that is, if the processing at all stages of perception is modulated by top-down conceptual influences—then there is no nonconceptual content and no philosophical argument could resurrect it. Furthermore, if there are perceptual mechanisms that retrieve information from a scene bottom-up then this information is the nonconceptual content of the states produced by these mechanisms. Thus, the existence of cognitively impenetrable mechanisms is both a necessary and sufficient condition for nonconceptual content. The task is to identify those contents of the perceptual cognitively impenetrable mechanisms that retrieve information that are also available to awareness. These contents constitute the phenomenal content of perception.

The paper consists of three parts. In the first part we present evidence from the neurosciences to support our claim that there is content retrieved from a scene through the mechanisms of vision in a conceptually unmediated bottom-up way. In the second part, we introduce the notion of ‘nonconceptual content’, which is defined as the content that is retrievable from a scene through the cognitively encapsulated processes of perception. Then we discuss the relation of this notion nonconceptual content with the notion of phenomenal content, as employed in the philosophical literature. In the third part, we specify the content in ‘nonconceptual content’. Finally, in concluding, we discuss some issues in the literature.

## **2. Evidence for Bottom-Up Representations of Objects**

### **2.1. The Evidence from the Cognitive Sciences**

We now present neurophysiological and neuropsychological evidence that supports the existence of a weak kind of object representation based on spatiotemporal, size, orientation, motion, and viewer-centered shape information that is retrievable bottom-up from a scene and that precedes other representations of the same object. There is a consensus that there are two (Goodale and Milner, 1992; Norman, 2002) or three (Glover, 2003) visual streams in the cortex, which serve roughly two different functions. The dorsal system utilizes visual information for guidance of action in one’s environment. For that purpose, it needs to have information about the dimensions of objects in body-centered terms. Thus, the information in the dorsal stream is transformed into an egocentric body-centered frame of reference, and uses viewer-centered structured representations of the surfaces of objects. The ventral system uses visual information transformed into a body-centered frame of reference for knowing one’s environment, that is, for identifying objects in it, for recognizing objects by comparing them with stored representations, and for storing new information in memory.

The dorsal system processes visual spatial information (information about the location of an object and of its relative position with respect to other objects), motion information, featural information regarding size, shape, orientation, and surface properties, and information regarding the affordances of objects in a body-centered frame of reference. This information is retrieved from the scene directly

by early vision, without recourse to any central higher processing (Glover, 2003; Goodale and Milner, 1992; Norman, 2002). The information is fed into the dorsal system through the magnocellular retinocortical pathways that carry spatiotemporal information, motion information, and information about size and shape. The information in the dorsal system allows a weak form of object representation that does not last more than a few seconds or minutes. Most of the mechanisms involved in the processing in the dorsal system operate outside of conscious awareness. Hence, the representations of objects in the dorsal system are not in general available to our conscious awareness (for exceptions see Norman 2002, pp. 89–90).

In contrast to the sparse information stored in the dorsal system, the representations of objects used by the ventral system are richer and constitute a strong object representation that contains semantic information. In addition to spatial information, motion information, and information about surface, size, orientation and shape, which is now cast in an object-centered reference framework, ventral representations contain information that requires top-down semantic influences deriving from specific object knowledge (information about fragility, temperature, function, weight, usage, functions, etc.). Indeed, the ventral system receives its information from the visual cortex through both the magnocellular retinocortical pathways and the parvocellular retinocortical pathways. The information requires reference to stored memories and relies on previously stored knowledge about specific objects. This representation is conceptually contaminated, since it is influenced by top-down conceptual inferences that reflect a large array of visual and cognitive interaction.

The above account suggests that spatio-temporal, shape, and size information processing precede processing of other featural information. This is supported behaviorally by data showing that on-line corrections are made very quickly (~100 msec) to changes in spatial attributes of targets. Furthermore, the magnocellular retinocortical pathways, which carry spatiotemporal information and information about shape and size and are projected to the superior parietal lobe of the dorsal system and inferotemporal cortex of the ventral system have fast responses. The parvocellular retinocortical pathways that project only to the latter and carry non-spatial information have slow responses. Neurophysiological and neuropsychological studies using *ERP* readings corroborate this showing that *ERP*s to the single-object vs. two-objects conditions diverge earlier than the *ERP* differences between target and non-target stimuli (Czigler and Balazs, 1998). This is evidence that the allocation of spatial attention to a particular object precedes even the categorization of the object's features as target related or non-target related or irrelevant.

Lamme (2003; 2004) offers a detailed account of the processes along the ventral pathway, and offers neural definitions of the psychological phenomena of perception, attention and awareness. He argues for two kinds of processing that take place in the brain, the feedforward sweep (*FFS*) and recurrent processes (*RP*). In the *FFS*, the signal is transmitted only from the lower (hierarchically) or peripheral (structurally) levels of the brain to the higher or more central ones. There is no

feedback; no signal can be transmitted top-down as in *RP*. Feedforward connections in conjunction with lateral modulation that occurs in the early perceptual areas extract high-level information that is sufficient to lead to some initial categorization of the visual scene and selective behavioral responses.

When a visual scene is being presented, the feedforward sweep reaches *V1* in about 40 ms. Multiple stimuli are all represented at this stage. At this level there are nonattentional selective mechanisms that prevent many stimuli from reaching awareness, even when attended to. Such stimuli are the high temporal and spatial frequencies, physical wavelength (instead of color), crowded or masked stimuli and so forth. The feedforward processing results in some initial feature detection. Then this information is fed forward to the extrastriate areas. When it reaches area *V4* recurrent processing occurs. Horizontal and recurrent processing allows interaction between the distributed information along the visual stream. At this stage, features start to bind and an initial coherent perceptual interpretation of the scene is provided. Initially, *RP* is limited to within visual areas; it is local. At this level one can be phenomenally aware of the content of perceptual states. At these intermediate levels there is already some competition between multiple stimuli, especially between close by stimuli. The receptive fields that get larger and larger going upstream in the visual cortical can not process all stimuli in full and crowding phenomena occur. Attentional selection intervenes to resolve this competition. Eventually, signals from higher cognitive centers and output areas intervene to modulate processing; this is global *RP*. At this level, one has access or report awareness of the information.

Lamme discusses the nature of information that has achieved local recurrent embedding. He suggests that local *RP* may be the neural correlate of binding or perceptual organization. However, it is not clear whether at this preattentive stage the binding problem has been solved. The binding of some features, such as its color and shape, may require attention, while other feature combinations are detected pre-attentively. So, before attention has been allocated, the percept consists of only tentatively but uniquely bound features that form the proto-objects (Vecera, 2000). Lamme (2003) and Lamme and Roelsfema (2000) argue that the  $2\frac{1}{2}D$  surface representation of objects and their surface properties are extracted during the local *RP* stage. Other research (Hildreth and Ulmann, 1995) suggests that spatial relations are extracted at this recurrent stage. In addition motion and size are represented in cortical areas in which local *RP* take place.

Since this information results from local *RP*, one can only be phenomenally aware of it. The information is situated between feedforward (unconscious) and globally recurrent (access conscious) processing. Thus, the locally recurrent processing is the neural correlate of phenomenal experience per se, or phenomenal awareness.

Within this framework, Lamme defines attention as a selection process where some input is processed faster, better and/or deeper than others do. Thus, it has a better chance of producing or influencing a behavioral response or of being memorized. Awareness is defined as the occurrence of recurrent processing

(*RP*). The processes in *FFS* are necessarily unconscious. When there is *RP*, awareness arises and so long as *RP* is limited (e.g. by attentional suppression) to early areas, there is only phenomenal awareness. Thus, we are phenomenally aware of the relevant states. When *RP* involves higher cognitive areas and output areas as well, then attentional selection intervenes to resolve competition and has an influence; it induces access awareness of the information extracted during that stage. Since the content of this awareness can be typically reported, this form of awareness is also called 'report awareness'. Hence, the type of information of which one is phenomenally aware is situated between feedforward (unconscious) and globally recurrent (access conscious) processing. According to Lamme (2003, p. 3) the information of which one is phenomenally aware is a short-lived, vulnerable, and not easily reportable form of visual experience. In contrast, information of which one has 'access awareness' is more stable and easily reportable.

We see that at the neural level attention and awareness can be defined as two fully separate mechanisms and thus constitute two different phenomena. Attention is the competition between neural inputs for output space. Awareness is the result of recurrent processing, independent of this competition, but its extent (and type) depends on this competition; attention determines when one goes from phenomenal to access awareness.

Lamme distinguishes between perception and awareness. Kanwisher (2001) makes the same distinction: one can be in perceptual states while being unaware that one is in such states. Perceptual awareness 'involves not only activation of the relevant perceptual properties, but the further construction of an organized representation in which these visual properties are attributed to their source in external objects and events' (Kanwisher 2001, p. 90). Thus, perceptual awareness presupposes binding of activated features with a representation that specifies a specific token, as opposed to type, object. This binding requires (2001, p. 108) 'visual attention', which thus becomes necessary for visual awareness of object tokens.

There are many ways one can be aware of a stimulus. One may be aware of a mere presence of certain of the features of the stimulus, of the category of the object present, or of the gist of a complex scene (Kanwisher 2001, p. 97). Thus, one can be visually aware of features or of the presence of an object and one can be also aware of the content of an organized representation. These are different kinds of awareness and only the latter requires visual attention. Thus, it is reasonable to argue that the distinction between phenomenal awareness and report awareness is latent in Kanwisher's work.

Kanwisher claims that the representations of objects that can be formed without attention (and thus, without report awareness) correspond to representations of types of objects rather than to representations of tokens of objects. This fits well with Lamme's account of vision where in the absence of attention, one can have only phenomenal awareness of some fleeting unstable representations of objects that contain sparse information and lack specific details about objects. It is plausible that these representations correspond to types rather than tokens of objects, types being at a more general level of abstraction and thus containing less information

than tokens. Thus, the contents of perceptual states of which one can be phenomenally aware correspond to types and not tokens of entities.

All these suggest that visual attentional mechanisms select features after these features have been detected and that even some kind of visual form may be detected before the onset of attention. There is also evidence that visual and even semantic features of objects can be detected even in the absence of attention. Studies of *Change Blindness*, *Inattentional Blindness*, and studies of perception in the absence of attention (Driver *et al.*, 2001; Kanwisher, 2001; Mack and Rock, 1998; Merikle *et al.*, 2001; Moore *et al.*, 1998) suggest that even when there is no awareness of stimuli when attention is diverted elsewhere, stimuli are nevertheless perceived, and grouping of features into some form of objects takes place along the ventral system. Moreover, there may be both perceptual and semantic processing of stimuli when the stimuli are not attended under conditions that preclude awareness of the stimuli. Thus, one should not hurry to assert that semantic top-down processing requires awareness, which justifies the choice to treat perception independent of awareness.

Further evidence regarding the precedence of representations encoding spatio-temporal, size, and form information, over representations involving conceptual top-down influences comes from positron emission topography (*PET*) and event-related potential (*ERP*) studies. Research with *PET* and *ERP* (Posner and Petersen, 1990; Posner and Raichle, 1994; Ziegler *et al.* 1997) literally provides a spatio-temporal picture 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 pseudo-words), (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 research sheds light on the role of top-down pathways.

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, pseudo-words 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, 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 pseudo-words 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 semantically analyzes the visual word forms. It should be noted 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.

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 appear. This means that the brain registers the word form 50 ms later than the visual features.

In other *ERP* studies, subjects were asked to search for a thickened letter in letter strings. This is 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. However, 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, since the same electrodes are activated, (mainly right posterior areas), in which the bottom-up registration of visual features occurs, with a time delay of about 100 ms. 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.

In other studies (Luck and Hillyard, 2000) with subjects who were instructed to attend to the left visual field in some trial blocks and to attend to the right visual field in other trial blocks, the subjects were asked to respond when they detect an infrequent target stimulus among the nontarget stimuli at the attended location. The *P1* wave (a component of the *ERP* waveforms) is larger in amplitude for stimuli presented at the attended location than for stimuli presented at the unattended location. Since the difference is due to the attended location, it is reasonable to assume that the amplitude of the *P1* wave is modulated by spatial attention. The effect begins 70 to 90 ms after stimulus onset; thus it is clearly an early perceptual and not a postperceptual effect. Spatial selective attention increases the activation of the neural sites tuned to the selected loci. The effect is sensitive to stimulus factors such as contrast and position and insensitive to cognitive factors. It occurs before the identification of the stimuli and is insensitive to the identity of the stimuli. It is

also independent of the task-relevance of the stimulus, since it is observed for both targets and nontargets.

The neural site of *P1* seems to be the *V4* area in the extrastriate cortex, which implies that spatial attention influences visual processing in extrastriate areas of the visual cortex. The initial sensory processing in the primary visual cortex *V1*, on the other hand, seems to be unaffected by attention (Luck and Hillyard, 2000). Thus, the *P1* component may represent the earliest stage of visual processing that is modulated by voluntary spatial attention (Mangun *et al.*, 2000). Attention seems to play two roles at the early stages of perceptual processing. It resolves ambiguous neural coding by suppressing competing input sources, and it improves signal-to-noise-ratios. Thus, the voluntary control of spatial attention is driven by stimuli and task-demands, not by cognitive demands.

Here is the overall picture. When a visual scene is being presented to the eyes, the feedforward sweep reaches *V1* at a latency of about 40 ms. Then this information is fed forward to the extrastriate areas, parietal, and temporal areas of the brain. By 70–80 ms after stimulus onset most visual areas are activated. The preattentive feedforward processing culminates within 100 or 120 ms after stimulus onset. 70 to 90 ms after the stimulus onset, spatial attention by modulating the *P1* waveform enhances visual processing in a voluntary task-driven search at the salient locations. 100 ms after the presentation of the stimuli at those locations 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 or structural description of the objects in a visual scene, and about 200 ms after it the voluntary task-driven search is registered in the same areas that process the visual features, which thus are modulated by object-centered attention. Thus, the top-down effects of attention to features are delayed in time, involve the same anatomical areas as passive perception, except that attention amplifies the recordings in these areas. Finally, about 250 ms after the stimulus, some of the same areas participate in the semantic processing of the input.

Attention intervenes in two stages to perform two different functions. First, in the form of spatial attention via the *P1* component in the early stages of perceptual processing to focus attention on the salient locations. Second, in the form of feature selection, and after some of the features retrieved in a bottom-up manner from the scene have fused to deliver the physical form of the objects, to enable binding of all features to form the objects of our experience and to solve the competition problem. These two attentional stages correspond to two different attentional circuits. One is the spatial attention that intervenes very early to orient attention specific places and the other is the 'visual attention' that intervenes much later and binds together features to objects perceived as being such and such in such and such a place and in such and such a time. The effects of visual attention are clearly postperceptual, and thus, do not affect the issue of the cognitive penetrability of perception. However, spatial attention seems to threaten the cognitive encapsulation of perception.

We said that the *P1* effect is insensitive to cognitive factors, and that this effect may represent the earliest stage of visual processing that is modulated by voluntary

spatial attention. This seems to imply that, through the modulation of attention, early visual processing is cognitively driven, since where to focus attention may be a cognitively driven decision. To delineate this issue we must say a few words about the top-down control of attention. Top-down control of attention is related to the working memory, since sites of that kind of attentional control underlie systems of working memory. Working memory stores information and performs executive control governing retrieval, encoding, and commands for the expression of attention (Baddeley, 1995). These two functions underlie the distinction in the attentional control processes between expectancy of an upcoming event and preparation for that event. The expectation of an event is not necessarily accompanied by an attentional preparation for it. The top-down attentional control of perception consists mainly in the attentional preparation for an upcoming event (Lagerge, 2000).

Thus, the claim that the *P1* component of spatial attention is insensitive to cognitive factors despite the fact that focusing attention may be cognitively-driven amounts to the following: cognitive factors do not control directly the *P1* effect. They determine the expectation for an event; but this is not sufficient to ensure attentional preparation for that event. Information regarding the upcoming display of an object may be kept in working memory while selective attention may be directed elsewhere. Selective attentional preparation for an event will follow the expectation of that event, if the event is task-relative. Once attentional preparation is effectuated the amplitude of the *P1* wave is only stimuli-driven.

Cognitive factors control the expectation of an event. Task-relevant factors 'translate' the expectation to attentional preparation for that event. Once the latter is on-line, neither cognitive nor task-relevant factors influence the *P1* effect and hence that stage of perceptual processing. Only stimuli factors do. It seems, thus, that once spatial attention has selected some loci for focusing information is registered at that stage of processing irrespective even of task demands. In other words, where to look may be determined in a top-down manner by cognitive factors. However, what one perceives at the relevant location, and hence the nonconceptual content of the perceptual states depends only on what is there.

We started by discussing evidence from neuropsychological and neurophysiological studies, which suggest that a representation of objects based on nonconceptual spatiotemporal, size and shape, information that is retrieved bottom-up from a scene is built in the dorsal system and the ventral system prior to a richer semantic representation built in the ventral system only. However, as with all empirical work, and this is true for the neurosciences with much greater force, data are never conclusive and constantly new data appear that either corroborate or refute older data. So, the conclusions drawn from the empirical studies discussed here are far from uncontroversial and far from being immune to revision. Our point, thus, is not that scientific inquiry has settled the issue once for all, but rather, than at present there is a substantial body of scientific evidence supporting the claims advanced in this paper. To support our claim further, we briefly review evidence from object-centered segmentation and attention studies.

## 2.2. Evidence from Object-Segmentation Studies

When we perceive a scene the first thing we usually do is to parse it in discrete objects, thereby individuating them. Our visual system does that by bundling parts of the visual field together as units. This function is called object segmentation and consists in a set of visual processes that determine which features combine to form the objects present in a visual scene. These processes segment a shape from the background and segregate it from other shapes that are similarly segmented from the background.

The preattentive perceptual grouping criteria that determine the regions of the image that are grouped together to determine proto-objects guide the distribution of visual attention, in the sense that spatially separated proto-objects in the visual field can be preattentively individuated and indexed and, as a result, the visual system acquires direct access to these objects. These perceptual grouping criteria are implemented in preattentive purely perceptual bottom-up segmentation processes (which means, no top-down cognitive effects modulate processing at this stage), which parse in a preliminary way a scene into objects. Vecera (2000) defines object-based attention as the visual processes that select a segregated shape from among several segregated shapes. Object segmentation takes place at many different levels of visual processing, both early and late, and a significant amount of object segmentation occurs in the preattentive stage of early vision (Driver *et al.*, 2001; Pylyshyn, 2001; Scholl and Leslie 1999; Scholl, 2001; Vecera, 2000).

The preattentive segmentation process results in the proto-objects. After proto-objects have been segmented, the visual system assigns to them indexes. The indexing may result in the visual system opening 'object-files' for the segmented objects in a scene. Any changes in features in the object file are construed as changes in the specific object; as a result the visual system sees an object undergoing changes and not a new object appearing. Thus, the role of object files is to maintain spatial and temporal continuity of objects in motion and change.

Research on object-centered segmentation and attention (Carey and Xu, 2001; Czigler and Balazs, 1998; Pylyshyn, 2001, 2003; Pylyshyn and Storm, 1988; Scholl and Leslie 1999; Scholl, 2001; Spelke, 1992) suggests that both adults and infants use spatiotemporal information first in order to individuate and track objects in a visual scene. Object individuation based on spatio-temporal criteria (for example, temporal synchrony or continuity and proximity) precedes and often overrides object individuation based on other featural criteria (for example shape and color). In other words, when a scene could be parsed, or objects could be individuated, in different ways depending on whether one uses spatiotemporal or other featural information, then the visual system of both adults and infants uses spatiotemporal information first: this is the principle of spatiotemporal priority. Only if spatiotemporal information fails to individuate objects are other features used to that effect.

This implies that representations of objects based on spatiotemporal information are constructed and used first, before representations that include size, shape, color, and other non sensible or semantic information, such as the kind of the object

involved. These representations are used to individuate, index, and track objects in a visual scene; that is, in picking up parts of the scene as things that persist in space and time and following them in space and time. In other words, there is a mechanism in vision that relates parts of visual representations to parts of the visual world.

Xu and Carey (1996) showed that 10-month-olds can employ spatiotemporal information to infer the existence of occluded objects behind a screen but cannot employ feature information for the same purpose. The objects are individuated by indexes that are feature-blind. 12-month-olds possess the capacity to use both kinds of information. Moreover, spatiotemporal criteria override conflicting feature information. Using two objects differing both in their perceptual properties and their categorical kind (a yellow duck and a white truck) they observed that infants are not surprised (look longer) when the truck disappears behind a screen and the duck appears, but they were surprised when nothing reappeared after entering the screen. This is evidence that a mechanism tracking the spatiotemporal history of objects is already in place, allowing the infant to individuate and follow the movement of objects, whereas, the feature tracking mechanism that identifies objects is overridden.

Spelke (1992) draws similar conclusions and argues that infants distinguish between featurally identical objects on spatiotemporal information and that spatiotemporal information is used for object individuation and numerical identity. Kahneman and colleagues (1992) show that features of individuated objects may change while the object is still seen as the same object as before. Tremoulet *et al.*, show (2000) that 12-month-olds use shapes and colors to individuate objects. However, this happens only when spatiotemporal information is ambiguous and the object segmentation processes fail to assign separate object-files to the objects in the scene. In this sense featural information may be used for such purposes whenever spatiotemporal information is ambiguous.

The research covers evidence from two domains: the development of object individuation and representation in infants and the theory of object-centered attention. It suggests that (a) both infant and adult individuation is based predominantly on spatiotemporal information; and (b) in both systems featural information is used for object individuation only if spatiotemporal information is ambiguous.

We contrast 'object identification' and 'object recognition' with 'object individuation'. 'Object identification' attempts to convey the notion of an object, which is represented as being such and such, that is, as falling under a certain description. This notion involves a strong conceptual component, in the sense that the object represented has been compared with other objects in some knowledge basis in memory and identified as being such and such. Object identification presupposes the existence of a concept associated with that object.

The term 'object individuation' involves a weaker kind of representation. It purports to convey the sense that an object file has been opened for that specific object, that the object has been 'catalogued' or 'indexed' as something that exists

and persists separately of other objects with its own continuous spatio-temporal history. In other words, ‘object individuation’ purports to convey the perception of the ‘objecthood’ of objects (the term ‘perception’ implies that there is no conceptual involvement in object individuation).

Research on object-centered segmentation suggests that the representation of objects as discrete entities that persist in space and time that is based on spatio-temporal information retrieved bottom-up from a scene precede a representation of the same objects based on featural information also retrieved bottom-up. These representations allow object individuation and both precede representations based on semantic information that allow object identification. The common thread of the research is the claim that there exists a level of visual processing in which objects present in a scene are parsed and tracked as distinct individual objects without being identified as particular objects. Object segmentation processes create object-files for the discrete objects parsed in a scene. The files individuate and index objects and are allocated and maintained primarily on spatiotemporal information. Since the information required for opening and maintaining the object files is retrieved bottom-up, the corresponding representational states have nonconceptual content.

### **3. Nonconceptual Content, Phenomenal Content, and Phenomenological Experience**

The terms ‘perception’, ‘sensation’ and ‘cognition’ are not always used with the same meaning, and since we will use them with their traditional meanings, it is better to clarify what we mean by each one of them. The distinctions are based on the distinction between early vision and high vision, and the distinction between feedforward processing and local and global recurrent processing (*RP*). Recall that these two distinctions overlap.

All processes that apply to the information contained in the retinal image, that is, processes that compute information on light intensity constitute *sensation*. Sensation includes processes that compute changes in light intensity by locating and coding individual intensity changes. It includes Marr’s *raw primal sketch* that provides information about zero crossings, bars, blobs, boundaries, edge segments etc. Since much of the information regarding surfaces is encoded in changes in the intensity of reflected light on the retina, the initial task of the visual system is to decode this information by locating, representing, and interpreting changes in intensity. Sharp intensity changes, for instance, are interpreted as surface boundaries. The properties of stimuli recorded at this level never reach awareness; there are non-attentional selection mechanisms involved here that filter out information. In neuroscientific terms, sensation consists in some of the early processes that belong to Lamme’s *FFS*. It might be that they occur before the binding of features extracted from the retinal image.

The processes that transform sensation to a representation that can be processed by cognition constitute *perception*. The output of these processes is a cognitively impenetrable content that is retrieved from a visual scene in a bottom-up way. A subset of this output, that which can be brought to phenomenal awareness is the 'phenomenal content'. In Lamme's theory phenomenal awareness requires local *RP*. It follows that only content that is formed by means of local *RP* can be "phenomenal content". As an example of perception, consider Marr's various grouping procedures applied to the edge fragments formed in the *raw primal sketch*. They yield the *full primal sketch*, in which larger structures with boundaries and regions are recovered. Through the *primal sketch* contours and textures in an image are captured in a bottom-up way. Although processing at that level involves lateral and local top-down flow of information, this, being within early vision, does not threaten the bottom-up character of the relevant processes. Perception comprises what is known as the intermediate-level vision that includes 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, since they rely on lateral and local top-down flow of information (Lamme and Spekreijse, 2000; Hildreth and Ullmann, 1995). In Marr's model the  $2\frac{1}{2}D$  sketch is the final product of perception. Spatial relations, position, orientation, motion, size,  $2\frac{1}{2}D$  shape, surface properties, and color are all bottom-retrievable in perception. In Lamme's framework, perception may consist of those stages of *FFS* that, along with lateral processing, bind features in the image, and of the stage of vision that involves local *RP*.

All subsequent visual processes fall within *cognition*, and include both the post-sensory/semantic interface at which the object recognition units intervene, as well as purely conceptual processes that lead to the identification and recognition of the array. This is the high-level vision or observation. In Marr's theory, the culmination of visual processes is the *3D model* of an object. The recovery of the objects cannot be purely data-driven, since what is regarded as an object depends on the subsequent usage of the information, and thus is cognitively penetrable. Several theories of vision hold that object identification is based on part decomposition that is the first stage in forming a structural description of an object and seems to depend upon knowledge of specific objects. Other theorists propose that objects are identified by template-matching processes (Edelman, 1999). 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. Similarly, template-matching relies on specific object knowledge and is, consequently, cognitively driven since the templates are the result of previous encounters with objects stored in memory.

We would like to explain here the qualification 'knowledge about specific objects'. Even if perception turns out to be of bottom-up character, it is still not isolated from knowledge. Knowledge intrudes on perception, since early vision is informed and constrained by some operational constraints that reduce indeterminacies in information (mainly the underdetermination of the  $2\frac{1}{2}D$  structure from

the 2D retinal stimulation). They are general assumptions about the world constraining visual processing (Marr, 1982; Spelke 1990; Pylyshyn, 1999; Ulmann, 1979). These constraints are not the result of explicit knowledge acquisition about specific objects but are reliable regularities about the optico-spatial properties of our world. This 'knowledge' is implicit, in that it is available only for the processing of the retinal image and cannot be overridden, whereas explicit knowledge is available for a wide range of cognitive applications and one can decide not to use it.

The distinction between a bottom-up, nonconceptual perception and a conceptual experience, puts us in a long tradition of similar distinctions. Jackendoff (1987) distinguishes 'visual awareness' from 'visual understanding', and considers Marr's 2½D sketch to be an exemplification of the former, and Marr's 3D sketch an exemplification of the latter. Similarly, Dretske (1993) distinguishes 'thing-awareness' from 'fact-awareness' and also (Dretske, 1997) a 'phenomenal sense of see' from a 'doxastic sense of see'. The first parts of the aforementioned pairs clearly correspond to a non-epistemic sense of perception. They involve phenomenal content. The second parts of the pairs correspond to an epistemic sense of perception. In the latter case, the content delivered is the content of judgments and beliefs.

Let us discuss the phenomenal content of perception and its relation to the nonconceptual content of perception. Phenomenal content, as is understood here, is a type of nonconceptual content; that nonconceptual representational content that is available to phenomenal awareness. Perceptual states have a phenomenal character in virtue of having some kind of representational content that makes things look in certain ways. The phenomenal content of experience includes representations of properties of things, events, or places, that these things, or events, or places have in virtue of appearing to us, or being disposed to appear to us, in certain ways. Suppose that *X* views *Y*. The phenomenal content of the perceptual states of *X* consists in representations of the properties of *Y* that *Y* has in virtue of appearing to *X*, or being disposed to appear to *X*, in certain ways. The phenomenal content is also called the content of sensation, sensuous content or sensation simpliciter. In all these denominations, 'sensation' is used as a generic term for visual perception. Since we have defined sensation in a radically different way, we retain the term 'phenomenal content' only.

The kind of content nonconceptual phenomenal states have depends, of course, on the definition of 'content'. We will say that a state has a representational content if it represents the world as being a certain way. In this case the state has a correctness condition, that is, the world can make the state come out true or false (Crane, 1992). Peacocke (1992) calls the correctness condition 'The minimal account of content possession'. It is the correctness condition and the possibility of falsehood or misrepresentation that render the nonconceptual content of perception intentional and thus representational, in the sense that it is about something in the world. According to Bermúdez (1995, pp. 344–349), the condition that a state should be semantically evaluable is a necessary but not a sufficient condition for the state to have representational content. For this reason, he adds

some further conditions that the content of a mental state must meet in order to be representational.

Tye (2002, p. 447) distinguishes between the conceptual content of a visual scene (that it represents, say, a tiger) and the phenomenal content of the same scene, which includes colors, shapes, and spatial relations obtaining among blobs of parts. He argues that the scene has the conceptual content it does, partly in virtue of its phenomenal content. Therefore, the conceptual content is grounded in the phenomenal content. Shoemaker (2002, p. 461) defines the objective properties of things as they appear in perception, as those properties that the subject perceives in Dretske's (1997) 'doxastic' sense of perceive, that is, those properties that the perceiver judges that things have when she takes her experience at face value. These objective properties are grounded in the ways things appear phenomenally, to wit, in the ways they look, in Dretske's 'phenomenal' sense of looking, to the perceiver. Similarly, Lowe (1992, p. 89) writes, 'such intentional content as a visual experience has must ultimately be grounded in its phenomenal or qualitative content'.

Thus, the phenomenal content plays a decisive role in the solution of the grounding problem. The reason is that a solution to the problem requires that the reference of the representational states with phenomenal content be determined not descriptively but causally. If phenomenal content should be determined by means of causal links with the world, then this content should be retrieved bottom-up from the visual scene through our visual system, without any conceptual involvement on the part of the perceiver; it should be nonconceptual. This is a natural link between 'phenomenal' and 'nonconceptual' content. Note, however, that our account of nonconceptual content does not rely exclusively on the need to ground phenomenal content. In that sense it is more liberal than other accounts that do.<sup>1</sup>

Peacocke (2001, pp. 240–241) claims that the ways properties or relations are given in experience are as important in characterizing the content of experience, as are the properties or relations that things are represented by the experience as possessing (see also Shoemaker 2002, p. 471). The ways things, events, or places seem to us are the phenomenal properties of experience. Peacocke (2001, p. 241) argues that the phenomenal content of experience represents things, events, or places and times in a certain way, as having certain properties or standing in certain relations, 'also given in a certain way'.

Tye (2005, pp. 18–19)<sup>2</sup> rejects this view. He claims that Peacocke models the awareness of qualities on the awareness of particulars. This is a mistake, because though individuals can look in various ways depending on the perspective of the perceiver, the qualities of which we are aware when we see these particulars do not

---

<sup>1</sup> We would like to thank an anonymous reviewer for raising this point.

<sup>2</sup> The references to Tye's (2005) are based on the pagination of the copy available on the internet in Tye's site.

look any way; they just show, as it were, their quality. 'Does it really make sense to talk of the way a shape is presented in experience or the way a color is presented? ... Redness, the property, is not experienced as being given in a certain way (other than as belonging to the thing).

Tye's view echoes Smith's (2003) claim that it is only perceptual experience, which makes us aware of objects, that can enjoy different perspectives (that make things look different ways to us), while sensations do not allow different perspectives. This is because perceptual experience presents objects as literally external to our bodies and this gives us the ability to examine perceptually various aspects of these objects from different perspectives. One cannot have a perspective of properties or qualities; it makes sense to say that I see a table that looks round from this view-angle, but that from that view-angle seems elliptical. This is not true with 'round'. If from another perspective one perceives 'elliptical', this is not the way the property 'round' looks from another perspective; it is a new property.

The way *X* can figure in the phenomenal content of an experience is the mode of presentation of *X* in that experience. On account of the fact that properties cannot be experienced as being given in a certain way 'our awareness of the relevant qualities is direct. It involves no mode of presentation' (Tye 2005, p. 18, fn. 20). Thus, it does not make sense to say that experience represents colors and shapes as given in a certain way.

As we have seen, to be in a perceptual state with a certain phenomenal content does not require that one possess particular cognitive capacities; hence the characterization of this content as nonconceptual. Crane (1992, p. 149) writes that *X* is in a state with nonconceptual content iff *X* does not have to possess the concepts that characterize the content in order to be in that state. Cussins (1990) also calls the content of such representations 'nonconceptual', which he defines as follows: for any state *S* with content, *S* has nonconceptual content *P*, iff the fact that *X* is in *S* does not imply that *X* possesses the concepts that canonically characterize *P*, meaning that *X* does not need to possess the concepts that would normally enter in a report of the content of *S* that adequately specifies that content. These definitions constitute the 'usual' way of understanding nonconceptual content.

As Tye (2005, p. 23) remarks, 'given the usual understanding of nonconceptual content, as far as the nature of the content itself goes, there need be no distinction between conceptual and nonconceptual content', as the difference between nonconceptual and conceptual content is a matter not of the properties of each kind of content but of whether one who possesses the relevant contentful states also possesses the corresponding concepts; your nonconceptual content may be my conceptual content, because you do not possess the required concepts. This leaves open the possibility that there be visual nonconceptual states having conceptual contents. Thus, the traditional definitions of nonconceptual content do not do a thorough job in distinguishing the two kinds of contents. They may lead someone to draw the wrong conclusion that nonconceptual content and conceptual content are on the whole the same (Crane 1992, p. 155). That conclusion is wrong because nonconceptual content and conceptual content are different kinds of content, not

contents whose sole difference lies in that conceptual content is accompanied by the possession of the relevant concepts. To put it differently, conceptual content is not simply nonconceptual content + concepts.

To overcome these difficulties, Tye (2005, p. 4) claims that 'visual experiences have contents that are robustly nonconceptual, and, insofar as they have such contents, they are nonconceptual states'. The robustly nonconceptual content is a possible state of affairs that contains entities of this world. The possible states of affairs consist of two basic types. First, structured complexes of specific particular items, properties, and relations. Second, structured existential states of affairs involving properties and relations and plausibly the subject of the experience. Nonconceptual content is a possible state of affairs of the second type. To explicate what he means, Tye provides the following example:

'I see the facing surface *S* of an object *O* and it looks red to me. My visual experience intuitively represents *S* as having the property of being red. At this level my experience is accurate if and only if *S* is red. But my experience has also something important in common with certain other visual experiences not directed at *S*. Suppose, for example, that *O* is replaced with another object *O'* that looks just like *O* or that I am hallucinating a red surface so that phenomenally it is for me just as it is in seeing *S*. Intuitively, in all three cases, it seems to me that *there is* a red surface before me. At this phenomenal level, my experience is accurate if and only if there is a red surface before me. This content is existential, not involving *O*.'

Tye seems to draw a distinction between visual experiences whose content represents a type 'a red surface' as opposed to a visual experience whose content represents a certain token 'the red surface *S*'. At the phenomenal level what matters is whether there is a red surface (a type) before me, not whether that surface is of *S* (a token). What both experiences with *O* and *O'* have in common is a red surface. Thus they are both instances of the same type of experience, namely 'here is a red surface before me'. This is entailed by (a) the fact that at the phenomenal level the correctness condition is the same both for the case in which I see *O* that has *S*, and in the case in which I see *O'* that does not have *S* (recall *S* is the surface of *O*); my experience is accurate iff there is a red surface before me, not if there is object *O* with surface *S* before me, and (b) the minimal account of content possession. Thus, at the phenomenal level only types can be represented, not tokens; phenomenal content is about types not tokens. From this follows that the nonconceptual phenomenal content is about types of entities not about entity tokens. It is not about structured complexes of specific particular items, etc. It is about structured existential states of affairs involving properties and relations.

We agree with Tye's notion of nonconceptual content. In the first part we argued that nonconceptual content represents types of objects and not object tokens. The representations of objects that can be formed in a bottom-up way without attention (and thus, constitute contents of which one can be only

phenomenally aware) correspond to representations of types of objects not of tokens of objects. Also, in the absence of attention, only some fleeting unstable representations of objects can be constructed that contain sparse information and lack specific details about objects. Since (a) the phenomenal content is nonconceptual content, (b) nonconceptual content is retrieved bottom-up before the onset of attention and (c) without attention the visual systems constructs representations of types and of tokens, it follows that phenomenal content represents types rather than tokens.

Tye's account of nonconceptual content involves no reference whatsoever to the conceptual framework of the perceiver, only to the kind of content of her states. A visual experience is robustly nonconceptual if it has the content described above. Though Tye does not address the issue of what determines this kind of content, we think he would not reject the view that the determinant of this content is the causal link of perceivers with the world.

Our account of perception is a causal theory of perception. For *A* to perceive *X* as being *F*, *A* must be related to *X* in such a way that it is necessary both that *X* is *F* and that if *X* were not *F*, then *A* would have had an experience with a different nonconceptual content. What ensures the right kind of causal relation between how things are and the way they are represented through the nonconceptual content of perceptual states is the fact that in perception information is retrieved bottom-up, and in that sense directly, from the environment. If the environment is thus then so it will be perceived and if it had been different the perceptual states it induces would have been different.

A reference to our causal links with the world in perception is also found in Peacocke (1992) and Stalnaker (2003). They maintain that the content of a state is nonconceptual if it is fixed externally to anything that the bearer of the state believes or knows (the externality criterion for nonconceptual content). This means that nonconceptual content is determined through causal links to the world, independent of what cognitive states the person, whose perceptual states have the content, is in. *X* is in a representational state with nonconceptual content *Y*, if *X* has (or is being disposed to have) a content that is causally (or nomologically) connected in a certain way to instantiated *Y*hood. Our thesis that nonconceptual content is retrieved from a scene in conceptually unmediated ways substantiates this, because information that is retrieved in conceptually unmediated ways is information that is retrieved bottom-up in a cognitively encapsulated way from a visual scene. This information is determined solely through some causal links with the world that are established by our visual system. The nonconceptual level delivered by one's perceptual systems is the causal product of the world's acting on our perceptual systems. We adopt, thus, a thesis similar to Fodor's (1998) 'semantic externalism', only in our case the causal links establish the meaning of the nonconceptual content and not the meaning of concepts.

Phenomenal content purports to capture the qualitative aspects of our conscious perceptual states. In this sense, it coincides with the sensations or qualia that accompany our experience and is interleaved with awareness. In the literature,

'sensation', 'phenomenal content', and 'awareness' are intermingled. When the term 'nonconceptual content' was introduced to capture an aspect of the phenomenal content, namely that to be in a state with phenomenal content does not require the possession or exercise of concepts, it became entangled in the web of awareness. Bermúdez (1995) saw that the studies of vision reveal processes that are unavailable to awareness and, yet, have representational nonconceptual content. This is the content of a class of certain computational subpersonal or subdoxastic states. Though, as we shall see, there are problems with Bermúdez' way of handling the nonconceptual character of these subpersonal states, the move is successful, since it opens the way of severing the link between awareness and nonconceptual content. Such a separation is made necessary by the fact that perception and awareness are separate phenomena.

Thus, it would be better to define nonconceptual content independently of awareness and then return back to discuss the relation between the nonconceptual content and awareness, and hence account for the phenomenal content. This is what we attempt to do in this paper. In the first part, we distinguished between perception and awareness and argued that perception can occur without awareness on the part of the perceiver. We introduced the notion of nonconceptual content, which is content retrieved bottom-up from a visual scene, independently of awareness and we said that a subset of the nonconceptual content, that which corresponds to states of 'perception' in which local recurrent processing occurs, is the phenomenal content of experience, as it is at that stage that the content of the processing states becomes available to phenomenal awareness.

The phenomenal content of the philosophers is accessible to introspection; thus, phenomenal content is accessible to consciousness. But which consciousness? We have seen that there are two forms of awareness. Phenomenal awareness and access or report awareness that are defined in terms of the kinds of processes that support them. Local recurrent processing that excludes memory and other cognitive areas of the brain, and global recurring processing involving memory brain areas, respectively. The two kinds of awareness have phenomenological differences as well. Things of which we are phenomenally aware have a shaky existence (in that they have a limited spatial and temporal coherence) that differs from the stable percept of which one has access or report awareness. We have a fleeting 'experience' of things of which we are phenomenally aware, since nonconceptual content, being about types and not tokens, lacks the coherence and specificity of conceptual content.

The phenomenal content, is accompanied by phenomenal awareness, which is distinguished from report or access awareness.<sup>3</sup> The difference between the two is not one of accessibility to conscious introspection or not; they both are. The difference is due to the role of attention. Phenomenal awareness does not require attention, whereas report awareness does. Report awareness is within the

---

<sup>3</sup> Block (1996) also distinguishes between phenomenal and access awareness.

conceptual realm (since it denotes contents that are reportable or contents to which we have access, and thus contents that could be conceptually articulated in principle), whereas phenomenal awareness is not. Thus, the two kinds of content are subject to different kinds of awareness and thus are contents of different kinds.

When we see a colored patch or a 3D shape we are consciously aware of that patch or shape. However, try to introspect some members of the set of the nonconceptual properties of perception, say a segment edge or the viewer-centered  $2\frac{1}{2}D$  shape of an object or the particular shade of a color. It takes mental effort to abstract from the attributes of objects as we experience them to introspect the segment edge or the viewer-centered  $2\frac{1}{2}D$  shape or that shade of a color. Even then, one has a 'fleeting' experience that lacks the sharpness of ordinary experience, on account of the fact that phenomenal contents lack spatial and temporal coherence. That is why one cannot retain a shade in one's memory for a long time, where as one can retain a 3D shape some times even for one's entire life span.

Usually we are not consciously aware of the phenomenal contents of perception, only of the experiential content; we are aware of a specific tiger, not of a type of a shape with a type of a color, though of course the tiger has a certain shape and color. The phenomenal content of perception is not necessarily accessible to report or access awareness. A particular shade of a color, for instance, cannot enter report awareness because it cannot be conceptualized, as far as its qualitative character is concerned, although one could conceptualize as does McDowell (1994) the specific experiential episode by uttering 'that shade'. When one has access awareness to phenomenal content, the content is necessarily embedded in one's conceptual framework and ceases to be nonconceptual.

We see now the root of the confusion that led philosophers to take phenomenal content to be of the same kind as conceptual content. When philosophers distinguish a nonconceptual 'phenomenal sense of see' from a conceptual 'doxastic sense of see' they assume that both are within the realm of consciousness, since the phenomenal content is thought to be that nonconceptual content of which we are aware (as opposed to the nonconceptual content of the subdoxastic states). Because of the lack of a distinction between phenomenal and access awareness, both phenomenal content and conceptual content were thought to be accompanied by the same kind of awareness. Thus, these senses of seeing seem to concern basically the same kind of content, whose only difference lies in the fact that the one but not the other is conceptualized. This led philosophers to lose sight of the fact that the awareness of a shade or of the surface of an object is not the same as the awareness of the shape on an object.

On our part, we retain the conceptual vs. nonconceptual distinction, but we do not subscribe to the latent thesis that both kinds of awareness are consciously accessible in the same way. The doxastic sense of seeing or report awareness is necessarily conscious, since it requires the employment of concepts. The phenomenal sense of seeing or phenomenal awareness, on the other hand, is not necessarily conscious, although its contents may be accessible to phenomenal awareness. Therefore, the term 'phenomenal' as used here is different from the term as

employed in the philosophical literature, on account first of the difference between phenomenal and access awareness, a distinction that lacks in the philosophical discussions of the issue, and second of the different mechanisms that subserve the two kinds of awareness (local vs. global *RP*).

Nonconceptual content has been defined in this paper as the content of those states that are formed during early vision. Nonconceptual representational content has been defined as the content of the states of perception. Nonconceptual representational content that is available to phenomenal awareness is the phenomenal content. In this view, conceptual and nonconceptual contents are the products of two different stages of vision; the former results from late vision or observation, where as the former from early vision or perception. The former requires global *RP* and thus involves both bottom-up processing and top-down effects from cognitive centers, whereas the latter requires *FFS* and local *RP* that involve bottom-up, lateral, and local top-down processes (the top-down processes are restricted within the system of early vision and thus do not infuse it with cognitive influences).

When the conceptual framework is being applied to visual processing, by means of the modulation of visual processing by top-down flow of information from cognitive centers to visual areas through the mediation of attention, the nonconceptual content does not simply get conceptualized. The conscious access to the content of experience permeates nonconceptual content with conceptual content. Because of the role of attention, the nonconceptual content acquires a coherence in space and time that it lacked before, its parts are combined in various ways to give rise to the rich percepts of our ordinary experience whose details stand in opposition to the dearth of information of nonconceptual states. Thus, conceptual objects are usually identified as being such and such, and our experience with previous instances of these or similar objects invests them with a host of properties that go far beyond what can be retrieved directly from a scene, such as membership to a category, its possible uses that go beyond its affordances, and so forth. Furthermore, nonconceptual content concerns types of entities not tokens of entities. In that sense it lacks the specificity of conceptual contents of experience. It is clear, in view of all these, that conceptual and nonconceptual contents have different properties and thus are different kinds of contents.

We agree with Heck (2001, pp. 514–515) that the content of experience of which we have report awareness and the ensuing beliefs that articulate it reflect or track but do not record the nonconceptual content of our perceptual states. The contents of the former, moreover, conceptualize the contents of the latter. This is where the grounding function of nonconceptual content lies; experiential content is systematically related and dependent upon the nonconceptual content of the perceptual states. Evans (1983, pp. 227–228) similarly claims that the conceptualization of nonconceptual content produces a cognitive state ‘whose content is systematically dependent upon the content of the informational state’.

As has been widely argued (Churchland, 1988; Dretske, 1997; Gregory, 1990; Hanson, 1958; Kuhn, 1968; Marr, 1982), the content of experience that identifies objects as being such and such has a strong conceptual ingredient. It is formed by means of both top-down and bottom-up processes, and thus has a strong conceptual component. One's intentional attitudes pervade experience, and its experiential content is the result of a cognitively penetrable process. The experiential content is already an interpretation of the input from the distal objects, an interpretation that is based on one's conceptual system. The act of perceiving involved is 'perceiving' in its doxastic conceptual sense. We call this kind of consciously accessible content the 'experiential' content of our experience.

In contrast to the experiential content, Pylyshyn (1999, p. 361) claims that the classes provided by the visual system are, at a first approximation, viewer-centered shape-classes expressible in the vocabulary of geometry, and are thus, different from the classes of objects into which we cut the world. He argues that the product of perception is the  $2\frac{1}{2}D$  sketch, which does not deliver objects but structured representations of surfaces of objects. The output of the encapsulated visual processor does not consist in the experience of objects but in a more abstract form of categories of such objects, classified according to their generic shapes. The objects that the phenomenal content of perception is about lie in the midway between the proximal stimuli and the full-blown objects of our cognitive lives. Pylyshyn (2001) calls such objects 'proto-objects'.

#### 4. What Is the Phenomenal Content of Experience?

We have said that viewer-centered shapes are standard items in any list of phenomenal content. Indeed, this analysis is justified as we have seen by the studies on vision. These studies show that some form of Marr's  $2\frac{1}{2}D$  models of the objects present in the scene, that is, representations of viewer-centered shapes are retrievable in a bottom-up way from the scene. The visual cognitively encapsulated processes that process spatiotemporal information, spatial relations, surface shading, orientation, color, binocular stereopsis, size, shape and movement involve local *RP*. This means that they are available to phenomenal awareness.

All the above are retrieved from the morphology of the scene by the mechanisms of early vision in a bottom-up manner. This presupposes a frame of reference within which the processes take place and which defines the center and orientation of the positional information. Since the mechanisms that extract this information are those of early vision, it is natural to assume that the resulting frame is viewer centered and that the viewer-centered frame is somewhat constructed on a frame provided by the human body.

Evans (1982, pp. 153–154) proposed the 'egocentric space' as an egocentric viewer-center frame of representing space and spatial properties. This space is defined by the three axes that center on the viewer's body, to wit the axes right-left, up-down, and in front-behind. Similarly, Campbell (1999, p. 86) argues that

there is a direct, non-observational knowledge of our own bodies, which he calls a 'body image'. This image gives the system a practical grasp of the ways it can act. It is characterized by an egocentric framework (defined the same way as Evans'), which is suitable both for place identification and guidance of simple actions (those, for instance, involved in 'indexical causality').

Peacocke (1992, pp. 61–62) provides exactly such an argument in elaborating his notion of scenarios. Peacocke searches for a level of nonconceptual content on which to anchor concepts in a non-circular way. This kind of nonconceptual content is provided by the spatial types 'the type being that under which fall precisely those ways of filling the space around the subject that are consistent with the correctness of the content'. To specify the spatial types one needs to fix an origin and the axes of the resulting frame. These elements cannot be defined with reference to the real world, since a spatial type may be instantiated at different places. Thus, the point of origin and the axes should be defined with respect to a thing that is always present irrespective of the location at which the type is instantiated, which is the body of the subject. The point of origin may be the center of gravity of the body and the axes of reference the directions of right-left, up-down, back-front, as defined on the subject's body.

Now, many different things can fill a spatial type. If these things are the surface and its orientation, texture, hue, brightness of light, degree of solidity of objects, as well as the direction, intensity and character of light sources (illumination conditions), and the rate of change of location, in relation to the origin and axes as defined above, then this type is a scenario. If one assigns time to a scenario and to its point of origin and axes real directions and places in the world, then one has a 'positioned scenario'. This positioned scenario is the representational nonconceptual content of the spatial type, and provides in part the phenomenal content. The content of a scenario is spatial representational content. Note here that the degree of solidity of objects probably is not directly retrievable from a scene. Its computation may involve comparison with objects stored in memory. In that sense, it relies on previously acquired knowledge, and thus, it requires a top-down flow of information from higher cognitive centers. Thus, it may not be part of the nonconceptual content of experience.

A scenario must be distinguished from any mental representation of this same content, although to describe a scenario one must employ concepts, in that the concepts employed in such a description are not themselves components of the representational content of the percept. In other words, one can have a certain type of perceptual experience without possessing the salient concepts that would describe it. Peacocke's scenarios, being types, can be instantiated by many different scenes.

To put things into the body-centered framework, however, does not suffice to ensure adequate descriptions of the contents of a scene for the purpose of actions. One also needs to consider the relative relations of the objects in the scene within the egocentric body image. Peacocke's (1992) proto-propositions provide such information. They constitute a layer of the nonconceptual content of perception.

One can, for instance, perceive a symmetrical relation without possessing the concept 'symmetrical'. Peacocke's (1992, p. 77) list of the contents of proto-propositions includes relations and properties, such as, 'symmetrical', 'same shape as', 'equidistant from', 'parallel to', 'square', 'curved', etc. Their content contains an individual, or individuals, together with a property or a relation. The proto-propositional content of a perception represents the property or relation as holding of the individuals it also contains. The relations may hold of either places or regions in space, or the objects located in those regions. Featural information is not encoded in proto-propositions, only spatiotemporal information. The content has the subject (the individual) predicate (the relation) form. This explains the 'proposition' in the 'proto-proposition'.

We have discussed evidence from the neurosciences that supports the claims made here concerning the priority and importance of a body-centered framework. The same evidence shows that spatial relations among objects, and the locations of objects are retrieved relatively early in vision and in a bottom-up manner. Thus, they are too, part of the nonconceptual concept of experience, their representational role consisting non only in the representation of spatial relations, but also in setting the whole frame in which the other kinds of nonconceptual representations take place.

The bottom-up processes of perception provide more than just object individuation and observable properties. They also seem to output functional properties of objects or affordances<sup>4</sup> and causal relations such as 'X "transfers" something to Y'. We will concentrate first on the issue of the causal relations. Petitot (1995) talks of the positional (local) content-structure of the scene. The positional content is nonconceptual, and conveys information about nonvisual properties, such as causal relations (e.g. X 'transfers' something to Y). Suppose that one witness a scene in which X gives Z to Y. The semantics of the scene consists of two parts: (i) the semantic lexical content of X, Z, Y and 'give' as a specific action, and (ii) the purely positional local content. The latter is the image scheme of the 'transfer' type. X, Y, and Z occupy a specific location in the space occupied by the scene (just as they are the arguments in the three-place predicate 'to give'). In the image scheme, X, Y, Z are thus reduced to featureless objects that occupy specific relative locations, and in that sense can be viewed as pure abstract places. More specifically, X, Y, and Z, which in a linguistic description of the scene are the semantic roles, 'are reduced to pure abstract places, locations that must be filled by "true" participants'. These places are related by means of an action, of a 'transfer' type (1995, pp. 251–252). Thus, between the structure of a scene as an organized whole and the structure of the sentence that describes it there is a homology. Petitot's places do not refer to the actual locations that are occupied by the objects in a scene. The whole

<sup>4</sup> Clark (1999, p. 346) defines 'affordance' as 'the possibilities for use, intervention and action which the physical world offers a given agent and are determined by the "fit" between the agent's physical structure, capacities and skills and the action-related properties of the environment itself'.

description can take place without any reference to the actual locations occupied by the objects. Petitot, in using the spatial metaphor, seems to allude to the notion of an object devoid of all features (including actual location) except that it persists in time individually retaining its identity.

The global positional information contained in a scheme is retrieved from the morphology of the scene by the mechanisms of early vision in a bottom-up manner. Petitot gives a detailed account of the routines and algorithms that might perform this task. One notes that the set of entities that fill Peacocke's scenarios is similar to the set of fundamental 'entities' which provide the elementary representational space.

Let us turn to the affordances of objects. Consider the interaction of a frog with a fly (Bickhard, 1996). The neural activity in the frog is not an internal representation of the fly. The representational content of the neural activity induced by the fly consists in the possibility of tongue flicking and eating on the basis of indications about potentialities that are afforded by specific objects. The content is about the potentialities or possibilities of further interactions that are afforded by the environment for the system's interactions with it. They implicitly predicate those interactive properties of the environment that could support its indicated interactions of the organism. They have a causal significance for the frog that does not rest in the internal representation of external things and in the possession of the notion of 'causality', but depends on the practical grasp by the frog of the implications of the afforded potentialities for its actions. Since this kind of content implicitly predicts the properties of the environment that afford actions it is representational but not conceptual.

The potentialities afforded by the environment for the system's actions have a causal significance only for that specific system and not for other systems, no matter how similar they are. This is so because the potentialities are tied with the disposition of the body and its specific position. Campbell (1999) calls this kind of causality 'indexical causality', because the content of the terms associated with it depends upon the context in which they are used. As the frog need not have an internal representation of the fly, so it does not need to know that the fly 'is within reach' to initiate action.

The notion of something being 'within reach' that we should want to use in characterizing its knowledge is dedicated to its own capacities for movement ... Its representation of something as 'within reach' may be quite directly tied to its own initiation of movement. The reason for saying that the representation is precisely a representation of something as 'within reach' is entirely its direct relation to the creature's actions (Campbell 1999, p. 85).

This kind of content is called 'construction-theoretic' by Cussins (1990). It allows an organism to navigate through, and interact dynamically with it.

Let us recapitulate our findings. The phenomenal content of experience consists in: (a) The individuated object that exists separately and has a spatiotemporal

history; (b) Some form of Marr's  $2\frac{1}{2}D$  representations of viewer-centered shape, size, orientation, and surface properties of objects; (c) The relational properties of the components of scene, whether it be places or objects, such as, symmetrical relations, topological relations, such as, up-down, left-right, etc; (d) The positional (local) content-structure of the scene, which conveys information regarding causal relation and the functional properties of objects.

## 5. Concluding Discussion

Our findings regarding the nature of nonconceptual content agree with most of the results obtained through the usage of the other methods we discussed in the introduction, but, there are also some striking discrepancies. The object-centered  $3D$  representations of objects that allow object identification are usually deemed to be part of the nonconceptual content. However, most theories of vision consider the  $3D$  representation of an object as the product of bottom-up nonconceptual and top-down conceptual flow of information. Peacocke also thinks that 'solidity' is nonconceptual, which we have reasons to doubt. Moreover, our list of nonconceptual content includes the functional properties that objects have in relation to some organism and some causal relations.

The notion of nonconceptual content defended here sheds light on a problem in Bermúdez' (1995), where he distinguishes the nonconceptual content of perception from the nonconceptual content of subpersonal or subdoxastic computational states. The difference between the two lies in the fact that one can be aware of the latter but not of the former. As an example of subpersonal states, Bermúdez mentions various computational stages of visual processes as presented in Marr's theory of vision, such as the zero-crossings. The content of perception is nonconceptual because one need not possess or exercise the requisite concepts in order for one's experience to have the phenomenal content that it does. The content of the subdoxastic states is nonconceptual because language does not have or cannot have the requisite concepts. As Bermúdez (1995, p. 353) claims, those subpersonal states that are candidates for having content (zero-crossings, for instance), have nonconceptual content due to the recondite nature of the requisite concepts.

We think that zero-crossings are not representational states at all, given that they are processes that compute changes in light intensity, thus belonging to sensation and not to perception. Suppose, however, that zero-crossings are representational states. Consider some future civilization that happens to have a language that included such concepts. Would that mean that for this civilization the content of zero-crossings would be conceptual? And what would that mean for the individuals in that civilization who are unaware of these concepts? And what would that make of the notion for our civilization that still does not possess the relevant concept? This is a clear case of one of the problems of the 'phenomenological' methods for determining nonconceptual content.

To account for why the content of computational states is nonconceptual, Bermúdez relies on the contingent fact that the subjects' language does not contain concept-terms that conceptualize the content. But this criterion is wrought with problems. Bermúdez must accept that if zero-crossings are nonconceptual just because language does not have the appropriate concepts, then for our imaginary civilization zero-crossings might become conceptual states. This does not mean of course that in order to perceive zero-crossings one has to possess the corresponding concept; it means that this content can be conceptualized. The way Bermúdez deals with this issue suggests that he construes nonconceptual content as conceptual content minus concepts. However nonconceptual content and conceptual content are different kinds of content. Even if one possesses the concepts that conceptualize zero-crossings, the kind of processing that results in states with zero crossings being cognitively encapsulated, these states have nonconceptual content.

Suppose, further, that no natural language could ever have such concepts because the content of the states whose content they are supposed to conceptualize is so fine-grained that it exceeds the limits of the expressive capabilities of any language. Then an independent non-circular account that does not involve the notion of nonconceptual content, must be given about how much fine-grained a content must be in order to exceed the expressive capabilities of a language; such an account is lacking.

In general, invoking awareness and availability of linguistic description renders Bermúdez analysis susceptible to two problems. First, in view of the fact that there exist two kinds of awareness, phenomenal awareness and access or report awareness, which kind of awareness befits the content of perceptual states? If it is phenomenal awareness, how is that different from report awareness? If it is report awareness, how is it that the content is nonconceptual despite the fact that report awareness involves global recurrent processing, which brings in the conceptual arsenal of the perceiver? Second, granting that perceptual content is nonconceptual because one can be in states with such content without possessing or exercising concepts, what is it that makes the content of subpersonal states nonconceptual? That we are not aware of such content is no proof that this content is nonconceptual. It could be that the processes that lead to the states with such content are modulated by our concepts (in which case the content would be conceptual, although not conceptualized), without us being aware of that modulation or without our language containing the relevant terms.

In our approach to nonconceptual content these problems vanish. The content of the states of zero-crossings is nonconceptual not because language does not or cannot contain the requisite concept-terms, but because that content is retrieved bottom-up from a visual scene. To justify the nonconceptual character of the contents of either phenomenal and subdoxastic stages, one need not take recourse to the resources of the language; it suffices to invoke the fact that these contents are retrieved from visual scenes in conceptually unmediated ways. Nonconceptual content is independent of awareness and independent of possession or exercise of

concepts. Thus, our account offers a unified explanation of why personal and subpersonal states have nonconceptual contents.

This account specifies that one could have only phenomenal awareness of nonconceptual content. Phenomenal awareness is different from the awareness that accompanies our experience of the world, which suggests that one should not hasten to draw conclusions regarding the nature of nonconceptual content based on introspection alone, because introspection usually taps on experiential content. It may be that one lacks concepts that would describe one's experience, but that does not mean that no concepts are involved in the formation of that content. One may think that one can see in the phenomenal sense things because one judges that one could have the content of the relevant states without exercising concepts, but in reality there is a cognitive modulation of the processes that led to the mental state with that content, as in the case of the 3D shape. This poses restrictions on the kind of awareness of nonconceptual content that stem from the nature of phenomenal awareness; one can be phenomenally aware of object-types not of tokens.

Tye (2002, pp. 454–455) claims that phenomenal content is poised, abstract, nonconceptual, and representational. We saw why it is nonconceptual, representational, and abstract. The phenomenal content is 'poised' because it is at the right level of processing to allow it to be fed into cognitive processing. The registration of differences in illuminances is a level too low, directly accessible to the cognitive centers. Our phenomenal content meets this requirement, since it is the outcome of perception. Given the way we have defined 'sensation', 'perception', and 'cognition', phenomenal content stands in the midway between, and connects sensation with cognition.

*Department of Psychology, University of Cyprus*

*Department of Philosophy and Social Sciences  
American College of Thessaloniki*

## References

- Baddeley, A. 1995: Working memory. In M. S. Gazzaniga (ed.), *The Cognitive Neurosciences* Cambridge, MA: The MIT Press.
- Ballard, D. H., Hayhoe, M. M., Pook, P. K. and Rajesh R. P. 1997: Deictic codes for the embodiment of cognition. *Behavioral and Brain Sciences*, 20, 723–767.
- Bermúdez, J. L. 1995: Nonconceptual content: from perceptual experience to subpersonal computational states. *Mind & Language*, 10 (4), 333–369.
- Bickhard, M. H. 1996: Troubles with computationalism. In W. O'Donohue and R.F. Kitchener (eds), *The Philosophy of Psychology*. London: Sage.
- Block, N. 1996: How can we find the neural correlate of consciousness? *Trends in Neuroscience*, 19, 456–459.

- Campbell, J. 1999: The Role of physical objects in spatial thinking, In N. Eilan, R. McCarthy, and B. Brewer (eds), *Spatial Representation: Problems in Philosophy and Psychology*. Oxford: Oxford University Press, 65–97.
- Carey, S. and Xu, F. 2001: Infant's knowledge of objects: beyond object files and object tracking. *Cognition*, 80, 179–213.
- Churchland, P. M. 1988: Perceptual plasticity and theoretical neutrality: A reply to Jerry Fodor. *Philosophy of Science*, 55, 167–187.
- Clark, A. 1999 An embodied cognitive science? *Trends in Cognitive Sciences*, 3(9), 345–351.
- Crane, T. 1992: The nonconceptual content of experience. In T. Crane (ed.), *The Contents of Experience: Essays on Perception*. Cambridge: Cambridge University Press.
- Cussins, A. 1990: The connectionist construction of concepts. In M. Boden (ed.), *The Philosophy of Artificial Intelligence*. Oxford: Oxford University Press.
- Czigler, I. and Balazs, L. 1998: Object-related attention: An event-related potential study. *Brain and Cognition*, 38, 113–124.
- Devitt, M. 1996: *Coming to our Senses: A Naturalistic Program for Semantic Localism*. Cambridge: Cambridge University Press.
- Dretske, F. 1993: Conscious Experience. *Mind*, 102 (406), 263–283,
- Dretske, F. 1995: *Naturalizing the Mind*. Cambridge, MA: The MIT University Press.
- Driver, J., David G., Russell C., Turatto M. and Freeman, E. (2001); Segmentation, attention and phenomenal visual objects. *Cognition*, 80, 61–95.
- Edelman, S. 1999: *Representation and Recognition in Vision*. Cambridge, MA: The MIT Press.
- Evans, G. 1982: *The Varieties of Reference*. Oxford: Oxford University Press.
- Fodor, J. A. 1998. *Concepts*. New York: Oxford University Press.
- Glover S. 2003: Separate visual representations in the planning and control of action. *Behavioral and Brain Sciences* (in press).
- Goodale, M. A. and Milner, D. A. 1992: Separate visual pathways for perception and action. *Trends in Neuroscience*, 15, 20–25.
- Gregory, R. L. 1990: *Eye and Brain: the Psychology of Seeing* (4th edn.). Oxford: Oxford University Press.
- Hanson, N. R. 1958: *Patterns of Discovery*. Cambridge: Cambridge University Press.
- Heck, Jr. R. 2000: Nonconceptual content and the 'space of reasons'. *Philosophical Review*, 109 (4), 483–523.
- Hildreth, E. C. and Ulmann, S. 1989: The computational study of vision. In Michael I. Posner (ed.), *Foundations of Cognitive Science*. Cambridge, MA: The MIT Press.
- Jackendoff, R. 1989: *Consciousness and the Computational Mind*. Cambridge, MA: The MIT University Press.
- Kahneman, D, Treisman, A. and Gibbs, B. J. 1992: The reviewing of the object files: object-specific integration of information. *Cognitive Psychology*, 24, 174–219.
- Kanwisher, N. 2001: Neural events and perceptual awareness. *Cognition*, 79, 89–113.
- Kuhn, T. 1968: *The Structure of Scientific Revolutions*. Chicago: Chicago University Press.

- Laberge, D. 2000: Networks of attention. In M. S. Gazzaniga (ed.), *The New Cognitive Neurosciences*, second edition. Cambridge, MA: The MIT Press.
- Lamme, V. 2004: Independent neural definitions of visual awareness and attention. In A. Raftopoulos (ed.), *The Cognitive Penetrability of Perception: An Interdisciplinary Approach*. Hauppauge, NJ: Nova Science Books.
- Lamme, V. 2003: Why visual attention and awareness are different? *Trends in Cognitive Science*, 7, 12–18.
- Lamme, A. F. and Spekreijse, H. 2000: Contextual modulation in primary visual cortex and scene perception. In M. S. Gazzaniga (ed.), *The New Cognitive Neurosciences*, second edition Cambridge, MA: The MIT Press.
- Lamme, V. A. F. and Roelfsema, P. R. 2000: The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neuroscience*, 23, 571–579.
- Lowe, E. J. 1992: Experience and its objects. In T. Crane (ed.), *The Contents of Experience: Essays on Perception*. Cambridge: Cambridge University Press.
- Luck, S. J. and Hillyard, S. A. 2000: The operation of selective attention at multiple stages of processing: Evidence from human and monkey electrophysiology. In M. S. Gazzaniga (ed.), *The New Cognitive Neurosciences*, second edition. Cambridge, MA: The MIT Press.
- Mack, A. & Rock, I. 1998: *Inattentional blindness*. Cambridge, MA: The MIT Press.
- Mangun, G. R., Amishi J. P., Hopfinger J. B. and Handy, C. T. 2000: The temporal dynamics and functional architecture of attentional processes in human extrastriate cortex. In M. S. Gazzaniga (ed.), *The New Cognitive Neurosciences*. Cambridge, MA: The MIT Press.
- Marr, D. 1982: *Vision*. San Francisco, CA: Freeman.
- Merikle, P. M., Smilek, D. and Eastwood, J. D. 2001: Perception without awareness: Perspectives from cognitive psychology. *Cognition*, 79, 115–134.
- Moore, C. M. and Egeth, H. 1998: Perception without attention: Evidence of grouping under conditions of inattention. *Journal of Experimental Psychology: Human Perception and Performance*, 23, 339–352.
- Needham, A. Baillargeon, R. 2000: Infants' use of featural and experiential information in segregating and individuating objects. *Cognition*, 74, 255–284.
- Norman, J. 2002: Two visual systems and two theories of perception: an attempt to reconcile the constructivist and ecological approaches. *Behavioral and Brain Sciences*, 25, 73–144.
- Peacocke, C. 1992: *A Study of Concepts*. Cambridge, MA: The MIT Press.
- Peacocke, C. 2001: Does perception have a nonconceptual content? *The Journal of Philosophy*, XCVIII (5), 239–269.
- Petitot, J. 1995: Morphodynamics and attractor syntax: constituency in visual perception and cognitive grammar. In R. F. Port and T. Van Gelder (eds), *Mind as Motion: Explorations in the Dynamics of Cognition*. Cambridge, MA: The MIT Press.
- Posner, M. I. and Petersen, S. E. 1990: The attention system of the human brain. *Annual Review of Neuroscience*, 13, 25–42.
- Posner, M. I. and Raichle, M. E. 1994: *Images of Mind*. New York: The American Scientific Library.

- Pylyshyn, Z. 1999: Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, 22, 341–423.
- Pylyshyn, Z. 2001: Visual indexes, preconceptual objects, and situated vision. *Cognition*, 80, 127–158.
- Pylyshyn, Z. and Storm, R. W. 1988: Tracking multiple independent targets: evidence for a parallel tracking mechanism. *Spatial Vision*, 3, 178–197.
- Scholl, B. J. 2001: Objects and attention: the state of the art. *Cognition*, 80, 1–46.
- Scholl, B. J. and Leslie, A. M. 1999: Explaining the infant's object concept: beyond the perception/cognition dichotomy. In E. Lepore and Z. Pylyshyn (eds.), *What is Cognitive Science?* Malden, MA: Blackwell.
- Smith, A. D. 2002: *The Problem of Perception*. Cambridge, MA: Harvard University Press.
- Shoemaker, S. 2002: Introspection and phenomenal character. In D. J. Chalmers (ed.), *Philosophy of Mind*. New York: Oxford University Press.
- Spelke, E. S. 1990: Principles of object perception. *Cognitive Science*, 14, 29–56.
- Stalnaker, R. 2003: What might nonconceptual content be? In York H. Gunther, (ed.), *Essays on Nonconceptual Content*. Cambridge, MA: The MIT Press.
- Treisman, A. 1993: The perception of features and objects. In A. Baddeley and L. Weiskrantz (ed), *Attention: Selection, Awareness, Control*. Oxford: Oxford University Press.
- Tye, M. 2002: Visual qualia and visual content revisited. In D. J. Chalmers (ed.), *Philosophy of Mind*. New York: Oxford University Press.
- Tye, M. 2005: Nonconceptual content, richness and fineness of grain. Forthcoming in T. Gendler and J. Hawthorne (eds), *Perceptual Experience*. Oxford: Oxford University Press.
- Ulmann, S. 1979: *The Interpretation of Visual Motion*. Cambridge, MA: The MIT Press.
- Vecera, P. 2000: Toward a biased competition account of object-based segmentation and attention. *Brain and Mind*, 1, 353–384.
- Xu, F. and Carey, S. 1996: Infant's metaphysics: the case of numerical identity. *Cognitive Psychology*, 30, 11–153.
- Ziegler, J. C., Besson, M., Jacobs, A. M., and Tatjana, A. N. 1997: Word, pseudo-word, and nonword processing: a multitask comparison using event-related brain potentials. *Journal of Cognitive Neuroscience*, 9 (6), 758–775.