Available online at www.sciencedirect.com

SCIENCE DIRECT®

Neural Networks 17 (2004) 861–872

Neural Networks

www.elsevier.com/locate/neunet

2004 Special Issue

Separate neural definitions of visual consciousness and visual attention; a case for phenomenal awareness

V.A.F. Lamme[a,b]

[a]Department of Psychology, University of Amsterdam, Room A626, Roeterstraat 15, 1018 WB Amsterdam, The Netherlands
[b]The Netherlands Opthalmic Research Institute

Received 17 February 2004; accepted 17 February 2004

Abstract

What is the relation between visual attention and visual awareness? It is difficult to imagine being aware of something without attending to it, and by some, visual consciousness is simply equated to what is in the focus of attention. However, findings from psychological as well as from neurophysiological experiments argue strongly against equating attention and visual consciousness. From these experiments clearly separate neural definitions of visual attention and visual consciousness emerge. In the model proposed here, visual attention is defined as a convolution of sensori-motor processing with memory. Consciousness, however, is generated by recurrent activity between cortical areas. The extent to which these recurrent interactions involve areas in executive or mnemonic space depends on attention and determines whether a conscious report is possible about the sensory experience, not whether the sensory experience is there. This way, a strong case can be made for a pure non-cognitive form of seeing, independent of attentional selection, called phenomenal awareness. This can be dissociated from the reportable form, depending on attention, called access awareness. The hypothesis explains why attention and consciousness seem so intricately related, even though they are fully separate phenomena.

Keywords: Visual attention; Awareness; Consciousness; Cortex; Recurrent processing; Feedforward sweep; Phenomenal experience; Iconic memory; Visual processing; Perception

---

1. Introduction

As soon as we open our eyes we have a rich experience of the scene that is in front of us. It is as if a picture of the outside world is generated in our head. Where is this picture coming from? How do nerve cells generate such an experience? These questions are at the heart of the search for the neural correlate of visual consciousness. Cognitive science is trying to unravel this mystery from two ends. By doing psychophysical experiments it is attempted to get a better grip on what visual consciousness actually is. What is this picture in our head? Is it really a full representation of the outside world, or is it largely an illusion? At the other end, neuroscience tries to establish what neural structures or processes are involved in generating this experience (Crick & Koch, 1998). The goal is to get an understanding of visual consciousness by a convergence of these two fields.

It seems like psychologists have to deal with the more difficult of the two problems. To know what this picture in our head actually is, they can only go by what subjects are saying or doing. They can ask subjects ‘what they see’, or do more complex analogs of that using psychophysical paradigms. In either case, however, they have to infer from behavioral measures what is within someone else’s mind. Some philosophers have argued that for that reason alone the search for the neural correlate of awareness can never be an ‘objective’ science. Others have countered this by stating that also such heterophenomenological observations are objective measures that can in principle be correlated with neural events (see Searle (1998) for both views). My viewpoint will be somewhere in the middle; I will try to argue that it is possible to know what someone else is seeing, but we should not simply take his word for it.

Combining these insights with many recent findings in the field of neuroscience enables us to get a clearer understanding of consciousness, and in particular its relation to and difference from visual attention. What I will try is to

---

E-mail address: v.a.f.lamme@uva.nl (V.A.F. Lamme).

0893-6080/$ - see front matter © 2004 Elsevier Ltd. All rights reserved.
doi:10.1016/j.neunet.2004.02.005

explain the very related phenomena of attention and consciousness at the basic neural level. It will be an endeavor in the spirit of cognitive neuroscience, where cognitive psychology and neuroscience meet, to obtain new definitions for behavioral and mental phenomena. I will mainly focus on trying to give such core definitions, mostly omitting or only referring to the experimental evidence supporting it, because that has been presented in earlier reviews already (Lamme, 2000, 2003; Lamme & Roelfsema, 2000; Lamme et al., 2000). Also, space does not allow me to relate these definitions to other, often very related, theories about the same issues, which is by no means to imply that what I write here is not inspired by what many others have produced.

2. The psychological/theoretical perspective

2.1. Awareness and attentive selection

Fundamental to the study of conscious experiences is the assumption that they are selective; we are not aware of everything we lay our eyes on. This is obvious from introspection, but even more dramatically demonstrated in so called change blindness (CB) and inattentional blindness (IB) experiments. CB occurs when subjects are viewing a scene, where one of the items changes position, color, identity, or simply disappears (Rensink, 2000, 2002; Simons & Levin, 1997; Simons, 2000a,b). Provided the image transients of such a change are masked, for example by interposing a brief blank interval between the two versions of the scene, subjects very often do not notice the change. CB occurs even when the change is as dramatic as the disappearance of a whole building, changing faces etc. Although most prominently demonstrated in natural or otherwise complex scenes, with many different objects, CB can also occur in relatively simple and abstract scenes, like the one shown in Fig. 1a, or even when there is only one stimulus in the display. Once noticed, or when pointed at it (see Fig. 1b), the changes are easy to detect, so CB is not a matter of low detectability. IB (Mack & Rock, 1998; Simons, 2000a,b) occurs when the subjects' attention is focused on a particular task, and stimuli are unexpectedly presented. When asked afterwards about these stimuli, subjects often cannot report about them.

Some have taken the CB and IB findings to imply that, even though we think we see everything that is in front of us, we actually have a very limited conscious representation of the outside world (O'Regan & Noe, 2001; Posner, 1994). At the least, the findings hint to a selective process, where a limited number of items reach a privileged status. Unless that state is reached, stimuli are not noticed.

What determines which objects reach such a conscious representation? Attention seems to play a crucial role here. Items that are attended, or grab attention by themselves, survive CB and IB (see Fig. 1b). Also, the number of items that may survive CB is approximately the same as the number of items that can be stored in working memory (which is four), and we know that storage into working memory depends on attention. It seems like attention guards the gate towards a representation that can be consciously reported or remembered (as in IB), or that can be compared with previous or succeeding stimuli (as in CB). We may summarize this view in a schematic as shown in Fig. 2a and e. Many sensory inputs reach the brain, and via the process of attentive selection, some of these reach a conscious state, which allows us to report about them.

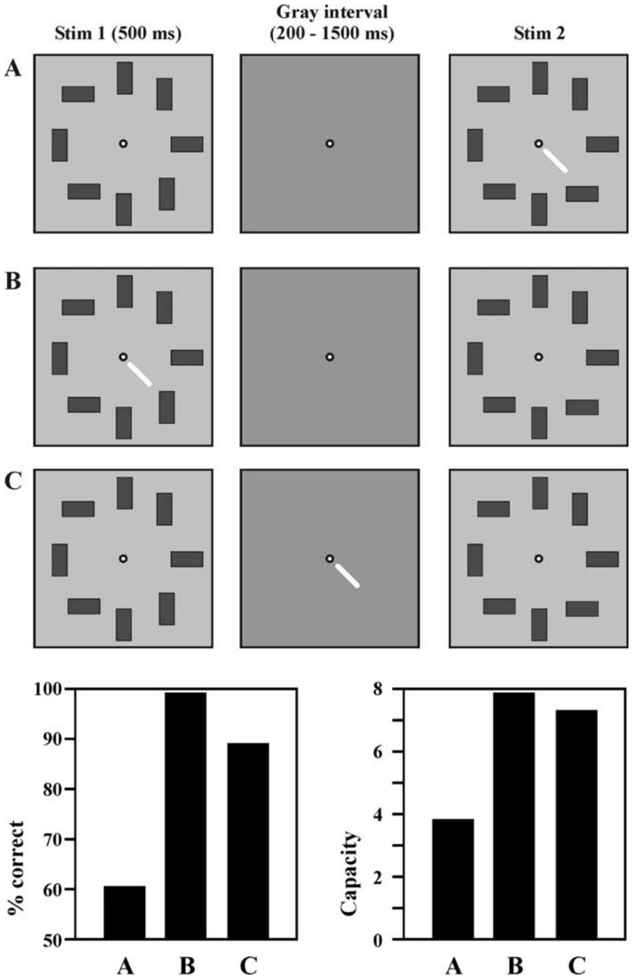

Fig. 1. Change blindness in an abstract scene, and the role of attention. In these change blindness trials (A–C), a scene containing multiple items is presented (stim 1), followed by a gray screen interval, after which the same scene (stim 2) is shown again. The subject is then asked whether the cued item (orange line) has changed or not (in this case it has; it changed orientation) (A). Subjects perform poorly at this task, (60% correct, lower left histogram). Performance can be converted in a 'capacity' measure, indicating (lower right histogram) how many items the subject had available (in working memory) for change detection. That number is about 4 items. When the to be changed item is cued in advance (B), subjects perform almost 100% correct (resulting in a virtual capacity of all 8 objects). However, when subjects are cued after the disappearance of stim 1, yet before the onset of stim 2 (C), subjects perform almost as well, and seem to have stored almost all objects.

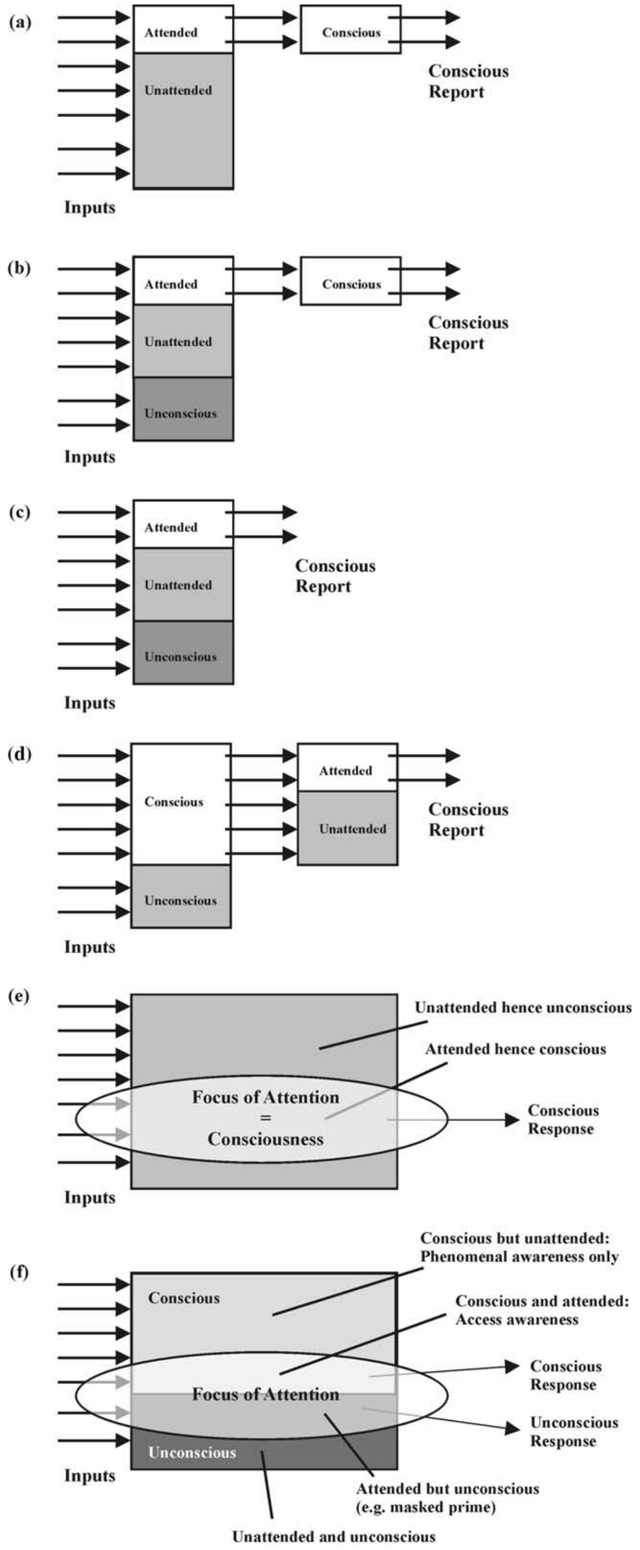

Fig. 2. Models of visual consciousness and its relation to attention. Visual awareness is limited, in the sense that we can only report about a small number of the inputs that reach us, typically those we attend to. It could be that attention determines what becomes conscious and what not, and hence determines what we can report about (a). However, there are also non-attentional selection mechanisms for consciousness (b). In these two views (a, b), there is no distinction between attention and consciousness, so that the latter term can be eliminated (c). Alternatively, the distinction between conscious and unconscious processing could be entirely separate from the attentional selection mechanism. In this view, many more inputs reach a conscious state, yet to report about these, we need attention (d). The serial models shown in a and d, are depicted in a more general way in e and f respectively. In e, consciousness is equated to what is inside the focus of attention. In f, the conscious/unconscious dichotomy is orthogonal to the attended/non-attended dichotomy. Hence, visual inputs can reach four different states: (1) conscious and attended, resulting in access awareness, (2) conscious yet unattended, leaving only phenomenal awareness, (3) unconscious but attended, which may result in the stimuli for which there is no phenomenal experience to still generate a response or influence behavior (as f.i. in masked priming), (4) unconscious and unattended. The fate of these inputs is uncertain.

2.2. Other forms of selection

Obviously, there are also properties of stimuli that may never reach consciousness, not even when attended. Many invisible stimuli or attributes activate neurons. Examples are high temporal and spatial frequencies (He & MacLeod, 2001), anti-correlated disparity (Cumming & Parker, 1997), physical wavelength (instead of color) (Zeki & Marini, 1998), crowded (Intriligator & Cavanagh, 2001) or masked stimuli (Enns & Di Lollo, 2000), or the non-dominant patterns during perceptual rivalry (Leopold & Logothetis, 1999). Also, fully attended stimuli may occasionally not be perceived, suggesting that sensory processing does not necessarily always complete to a perceptual stage (Super et al., 2001).

These are all non-attentional selection mechanisms. What is perceived during rivalry, for example, is hardly under voluntary control. These mechanisms likewise filter information before it reaches awareness. We could add these to the scheme of Fig. 2a, so that we have three classes or levels of processing that a visual input may reach: unconscious, unattended, or attended. Only the latter class reaches awareness (Fig. 2b).

There is something obviously awkward about this classification. There is no difference between attended and conscious stimuli, and as the neural basis of attention is fairly well studied (Desimone & Duncan, 1995) we would better eliminate one of the two terms (Fig. 2c). Some have indeed explicitly argued for this, in stating that in fact there is no awareness beyond attention (O'Regan & Noe, 2001). In that case, however, we would still need the term unconscious, for inputs that do not reach awareness even when attended.

2.3. Separating awareness from attention

A theoretically more logical solution is given in Fig. 2d. This model makes an early distinction between conscious and unconscious inputs, while the attentive selection process operates at an independent stage; attention does not determine whether stimuli reach a conscious state, but determines whether a (conscious) report about stimuli is possible. Likewise, attention determines whether items are stored in a sufficiently stable manner (working memory) to allow a report at a later time or to allow a comparison with a subsequent scene. Thus, the model would equally well support the CB and IB results; without attention, stimuli (or their change) cannot be reported. More precisely, the interpretation of the CB and IB experiments according to the model would be that we *do* have a conscious experience of many items in a scene. However, without attention these items are not stored in a sufficiently stable manner to allow a report at a later time (IB) or to allow a comparison with a subsequent scene (CB). In other words, attention is required to store items in working memory, and only those items survive CB or IB. The conscious experience we have of unattended items is very vulnerable, and apparently is overwritten as soon as a new scene hits the eyes, explaining CB. In other words, CB and IB are not necessarily failures of consciousness, but of conscious memory (see Wolfe (1999) for a similar argument).

Fig. 1 describes a psychophysical experiment in support of this view. We have seen that cueing the item that might change in a display of many objects protects from CB (Fig. 1b). Surprisingly, however, cueing the relevant item long after the first stimulus has disappeared, yet before onset of the second stimulus, also protects from CB (Fig. 1c) (Becker et al., 2000; Landman et al., 2003). Apparently, after the first display has disappeared, a neural representation of almost the whole scene is still present, and attention can select from this representation to store the relevant item in working memory. After the onset of stimulus 2, this representation has vanished, as cueing at that time does not help (Fig. 1a).

The model thus argues for the existence of a short-lived, vulnerable, and not easily reportable form of visual experience, which contrasts with a more stable, reportable form of awareness. A very similar distinction has been made by Block, who distinguishes between ‘phenomenal’ and ‘access’ awareness (Block, 1996). In the domain of sensory memory, a comparable distinction is made between a retinotopic, fleeting form (iconic memory; Coltheart, 1980) and a more durable non-retinotopic form (working memory; Levy & Goldman-Rakic, 2000). According to this view, attentional selection is inherently independent of either awareness or memory, yet determines whether we go from phenomenal to access awareness or from iconic to working memory.

An important feature of this model is that visual awareness is not simply equated to a conscious report. It is recognized that some selection process comes in between the conscious experience and the report about this experience. In some sense this could also be considered a decision process. In many psychophysical experiments, subjects have to say, in one way or another, either ‘yes’ or ‘no’ to the question ‘did you see the stimulus?’ Signal detection theory has shown us that these answers not only depend on the (subjective) experience of the stimulus (Wickens, 2002). Depending on the inclination to say either yes or no (the decision criterion) percentages yes and no may vary enormously for identically visible stimuli. In SDT, stimulus visibility (d') is dissociated from percentage correct via a mathematical model of the decision process. Similarly, the model of Fig. 2d argues not to equate awareness to the report about awareness. In this scheme, a conscious report is taken exactly for what it is, a motor output, and it is recognized that a decision process sits between the sensory experience and the motor output. For one thing, this approach will guard us from finding the neural correlate of awareness in the alpha motor neurons of the spinal cord.

Things find a very natural position in this scheme, which will prove very beneficial in understanding visual consciousness in neural terms, as I will elaborate on below. We either have a conscious experience about visual stimuli or not. Attention is a separate selection process, which is in principle independent of the conscious phenomenal experience (Fig. 2f). Attention is a limited capacity, bottleneck-like, process, that allows stimuli to be processed deeper or faster, and which is necessary for storage in a durable working memory store or for a conscious report about stimuli.

2.4. How to study this ‘hidden’ stage

Now that I have dissociated visual consciousness from a report about it, an immediate problem arises. How can we know what a subject is seeing if we cannot simply rely on his report about it? This is not as hopeless at it may seem. But let me begin by stating what one should not do in studying sensory experience. First of all, avoid attention and working memory being the variable. IB is sometimes taken to imply that there is no phenomenal experience of those stimuli. It can be successfully argued, however, that what is not present is conscious memory of those stimuli, and that a more appropriate term for the phenomenon would be inattentional amnesia (Wolfe, 1999). This would be fully in line with the model of Fig. 2d and f. Similarly, in CB experiments, what one is doing in fact, is asking a subject what he saw one image before the present one (or what he saw one eye movement ago). This implies we are asking the subject about what is in working memory. Obviously, if we want to know about sensory experience per se, it is better to ask the subject what he is seeing *now*. Iconic or sensory memory in that sense is a much better reflection of visual awareness than other forms of memory. Arguing that CB is evidence for a limited sensory experience is in fact the same as arguing that when someone has forgotten what he saw yesterday he was blind that day.

A second conclusion would be not to have decision processing interfere too much between phenomenal experience and a report about it. In that sense, the use of threshold like stimuli, such as in masking experiments, is to be avoided. This will reveal more about the nature of the decision process than about phenomenal experience.

When easily visible stimuli are used, it can be assumed that when they reach consciousness, and when the attention/report apparatus is properly allocated, a report will follow because its signal will always be above decision criterion. Better still, incorporate the decision process in the model that is used to analyze the behavioral data. By deliberately manipulating the decision criterion, it is possible to tease apart the decision process and the conscious/not conscious dichotomy (see Lamme, 2000; Super et al., 2001).

In summary, the conscious/unconscious dichotomy, even though not directly linked to a subjects' report, can be studied when the influence of any attentive selection or decision mechanism following this dichotomy is under full control, either by leaving it constant or by deliberately manipulating it. Second, it is important to come as close as possible to what the subject is seeing now. As we cannot really ask a subject what he sees at the moment of his response, iconic memory is probably the closest to actual phenomenal experience. The experiment of Fig. 1 makes the distinction between the two stages of awareness very explicit.

2.5. What is the difference between conscious and unconscious?

Having dissociated awareness from a report about it, another problem emerges. What exactly should we think of the distinction between conscious and unconscious in Fig. 2d and f if it is not strictly related to our own experience of reportability? From a definition point, the distinction is simple: unconscious stimuli or stimulus properties are those that we cannot report about, even when attended to. At first sight, this seems to result in a rather moot distinction between conscious and unconscious, being somewhat like 'we are not conscious of what is behind our back', i.e. of what the senses do not transmit. There are many more stimuli and stimulus properties, however, that we do not see, even though they evoke neural activity, not only in the eye and subcortical structures, but also in cortical areas (see Section 2.2). The distinction can be made much clearer, however, after we have considered the neural bases of the processes described thus far. I will return to the distinction between conscious and unconscious after that, in Section 4.

3. The neuroscience perspective

3.1. Starting points: processing and memory

In the cognitive neuroscience approach it is attempted to come to a better understanding of visual consciousness by converging insights from psychology and neuroscience. We therefore have to formulate specific ideas about the neural basis of psychological processes, ultimately not shying away from redefinitions of those processes. This has proven to be a very difficult problem. A universally accepted understanding of even the most elemental visual processes, such as motion perception, color constancy, or perceptual grouping, is still beyond us. So everything that is said about the neural basis of visual awareness is still very much in the hypothesis stage. Some even argue that neuroscience will never explain awareness. Therefore, I think it is good to give some starting points, about neural processes that we do understand, and go from there towards explaining complex issues like awareness and attention.

Although many issues in neural processing are not fully understood, we do understand them at a more principal level. We know that the senses transduce physical information from the outside into neural activity. We know that neurons integrate inputs at their dendrites, resulting in an output at the axon. These neurons are embedded in an intricate network in which we can identify nuclei, areas, pathways, modules, etc. In essence this is the basis of what we call sensory processing (leaving chemical neuromodulation out of the picture for the moment). The immense complexity of the anatomical connections between neurons renders the true nature of how successive neurons transfer information rather difficult to study. For that reason, many details, and also many fundamental properties, of how a sensory input is translated into a motor output still evade us. But still, we can in principle imagine how the brain might do that. If we start from our understanding of the reflex-arch, it is imaginable how we can build sophisticated input-output mappings from this principle (Colby & Duhamel, 1996). We may have to include very complex computational concepts, such as parallel pathways (Bullier, 2001), recurrent processes (Lamme et al., 1998b), synchrony (Engel et al., 2001), modulatory influences (Lamme & Spekreijse, 2000; Albright & Stoner, 2002), etc., of which we still understand very little, but there is no real explanatory gap there, only a lot of work still to do (comparable to, say, the field of genomics). In other words, sensory processing and sensorimotor transfer is understood at the fundamental level, although many important issues still need to be worked out.

A similar reasoning can be applied to memory. The synapses that mediate processing are plastic, so that the transfer of information may be modified. In this way, preceding events may induce changes to the network, and this is what we call memory. Again, the issue is highly complex, both with respect to the 'preceding events' as with respect to the 'changes in the network'. In principle, any preceding event that results in neural processing will have an influence on the network. Ongoing activity changes, some activity starts to reverberate in the network, and we enter the domain of working memory (Goldman-Rakic, 1996). Some (possibly LTP like) events may induce more lasting changes in synaptic transfer, which eventually may result in an anatomical consolidation of the changes in synaptic transfer (Dudai, 2002). Now we are in the domain of long- term memory. This is not to say that all aspects of memory and its neural basis are understood, but again we have at least some idea of basic principles.

3.2. Attention = processing × memory

Combining the core concepts of sensory processing and memory may be sufficient to explain visual attention (Desimone, 1996, 1998). Attention is a selection process where some inputs are processed faster, better or deeper than others, so that they have a better chance of producing or influencing a behavioral response or of being memorized. Attention induces increased (Desimone & Duncan, 1995) and synchronous (Fries et al., 2001) neuronal activity of those neurons processing the attended stimuli, and increased activity in parietal and frontal regions of the brain (Driver, 2001). The increased neural activity is in principle sufficient to explain why the associated stimuli are processed faster, deeper, etc. The main problem lies in explaining what brings the enhanced activity about.

Attention may be grabbed externally (Egeth & Yantis, 1997). Some stimuli are simply processed more efficiently than others. These stimuli we call salient. A bright stimulus will win from a dark one, a moving from a stationary, a foveal from a peripheral, etc. This is mainly due to the properties of the adult processing network, shaped by genetics and visual experience. In other words, saliency reflects how long-term memory has shaped and modified sensory processing.

But preceding stimuli may subtly change these properties. Imagine a stimulus entering the system, and another stimulus following within 100 ms or so. If the two stimuli share properties (such as retinal position), it is understandable why processing of the second stimulus will be more efficient than of a similar stimulus not preceded by the first stimulus. The first stimulus will have ‘paved the way’ in the sense that neurons are already activated above threshold, and this activity may persevere for some time (Fig. 3). This is a typical attentional priming situation (Dehaene et al., 1998). More specifically, processing of the first stimulus has led to a short-term memory trace (in this example maybe better called a sensory or iconic memory trace), and processing of the second stimulus is influenced

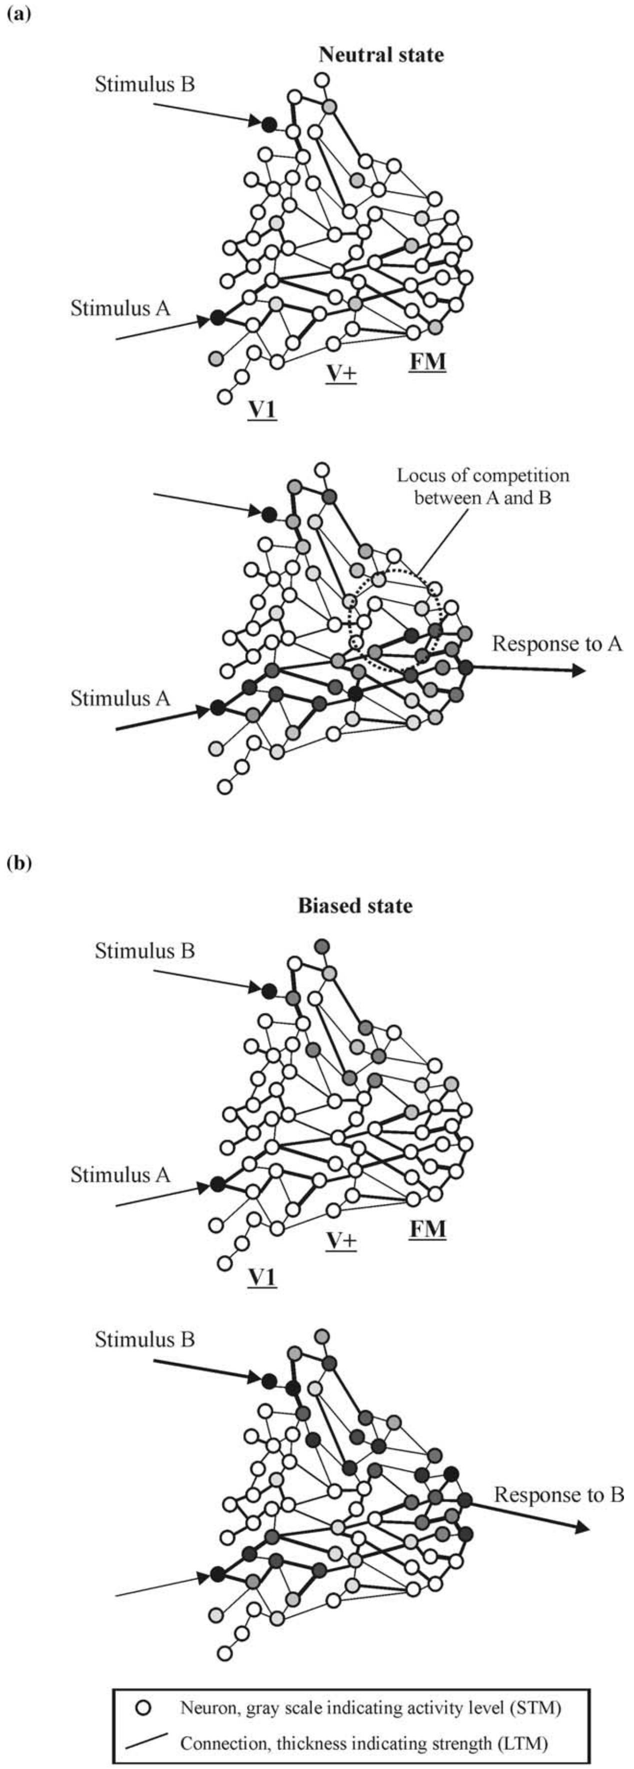

Fig. 3. Attentional selection is a convolution of memory and processing. Selection is necessary when two stimuli (A, B) reach the brain, yet only one response is possible. Competition, typically at the level of the extrastriate areas, prevents all inputs to reach output areas of the brain. Depending on the state of the brain when stimuli arrive, either of the two outputs may be selected: (a) shows a neutral state, where stim A is processed more efficiently, i.e. better matches stored synaptic weights: stim A is more salient, and the associated neural activity is stronger or more synchronous (darker dots); (b) shows a biased state, where the processing of a previous stimulus has left a short term trace of activity (light gray dots). Now, processing of stim B towards a response is favoured. Thus, attentional selection results from the convolution of the processing of current inputs with long and short term memory.

by this trace. Also inhibitory influences from the first stimulus are possible, for example at other locations, or at later times, when neural activity rebounds, resulting in inhibition of return (Egeth & Yantis, 1997).

With endogenous attention (Egeth & Yantis, 1997; Posner, 1994), the situation becomes more complex, but not fundamentally different. Now, an external event, such as an abstract cue, has to be translated in something akin to the 'paving of the way' described above. Parts of the brain that extract the meaning of the cue, and that are able to relate this to current needs and goals, must pre-activate or otherwise facilitate the appropriate sensory pathways, mostly via cortico-cortical feedback or subcortical routes. Regions that are able to do so will be at the interface between sensory and motor representations (parietal cortex), or will be where sensory, motivational and internal milieu information meets (pre-frontal cortex) (Driver & Frackowiak, 2001). Such top-down paving requires more time, yet it is more flexible and under voluntary control than bottom-up types of attention.

This is not to say, however, that something like intention or free will has to be incorporated in the idea. What we may experience as free will or intention, in this simplified scheme is nothing more than a combination of current and past inputs that operate on the current state of the network. We do not need anything else than the combination of sensory processing, the processing of internal milieu variables, and short and long-term memory to explain why a particular brain at a particular moment in time is inclined to favor one stimulus over another (Desimone, 1996).

So in summary, I think that all forms of attentional selection can be explained at the fundamental level as a convolution of sensory processing with short and long-term memory, even though, again, many details still need to be worked out. Therefore, I strongly argue that attentional selection is not a-priori associated with (visual) consciousness. We can imagine all the operations described above to occur in brains (or machines for that matter) without any phenomenal experience. We do not need to have an explanation for phenomenal experience to understand attention.

3.3. Visual consciousness = recurrent processing

What remains to be found, then, is a similar core understanding of phenomenal experience. We know that neural (including cortical) activation does not necessarily lead to consciousness. Hence the search for the Neural Correlate of Consciousness (NCC), where it is investigated what kind of neural activity is—and what kind is not—capable of producing awareness (Crick & Koch, 1998). Elsewhere (Lamme et al., 2000), I have argued that a strictly localizationist approach in the search for the visual NCC will be barren; there is no region in the brain whose activation automatically leads to visual awareness. The NCC is not anatomically defined, but functionally; some type of neural activity leads to awareness, while other types do not. With respect to that question, I have made a strong point (Lamme, 2000; Lamme & Roelfsema, 2000) of distinguishing between the so-called feedforward sweep (FFS) and recurrent processing (RP) (Lamme et al., 1998a).

The FFS is defined as the earliest activation of cells in successive areas of the cortical hierarchy. Typically, V1 starts to respond 40 ms after stimulus onset, and higher, extrastriate areas respond at successive slightly increasing latencies. At about 80 ms most visual areas are activated, at 120 ms visual activation can be found in all cortical areas, including motor cortex (Lamme & Roelfsema, 2000). Surprisingly, these early responses already fully express the receptive field (RF) tuning properties of cells, even complex ones like face selectivity in area IT (Oram & Perrett, 1992). Feedforward connections are apparently capable of generating sophisticated RF tuning properties and thus extracting high-level information, which could lead to categorization (VanRullen & Thorpe, 2001) and selective behavioral responses.

As soon as the FFS has reached an area, recurrent interactions between neurons within that area and neurons that have been activated earlier at lower levels may start (Fig. 4). These interactions are mediated by horizontal connections and feedback-feedforward circuits between and within areas (Lamme et al., 1998a). They are expressed as modulatory influences from beyond the classical, feedforward, RF (Albright & Stoner, 2002; Lamme & Spekreijse, 2000).

The hypothesis I put forward is that the feedforward activation of whatever area in the brain is not sufficient for visual awareness. Even when high level areas in temporal, parietal or frontal cortex are reached, this in itself does not lead to any phenomenal visual experience, i.e. is unconscious. Recurrent interactions between areas, most notably between V1 and extrastriate areas, are necessary for the visual input to reach consciousness. Some important observations can be made about the relation between FFS, RP, and visual awareness in support of that idea:

1. Backward masking renders a visual stimulus invisible by presenting a second stimulus shortly (f.i. 40 ms) after the first (Enns & Di Lollo, 2000). The masked stimulus, even though invisible, still evokes selective feedforward activation in visual and non-visual areas as widespread as V1, IT, FEF, and Motor cortex. Neurophysiological manifestations of recurrent interactions are, however, suppressed by backward masking (Lamme & Roelfsema, 2000; Lamme et al., 2002).
2. A stimulus becomes invisble when the features that define a shape from its background are switched in the two eyes, so that for example a red face on a green background in the left eye is combined with a green face on a red background in the right eye. Despite their invisibility, these stimuli activate the same cortical areas as visible versions (where the two eyes see the same

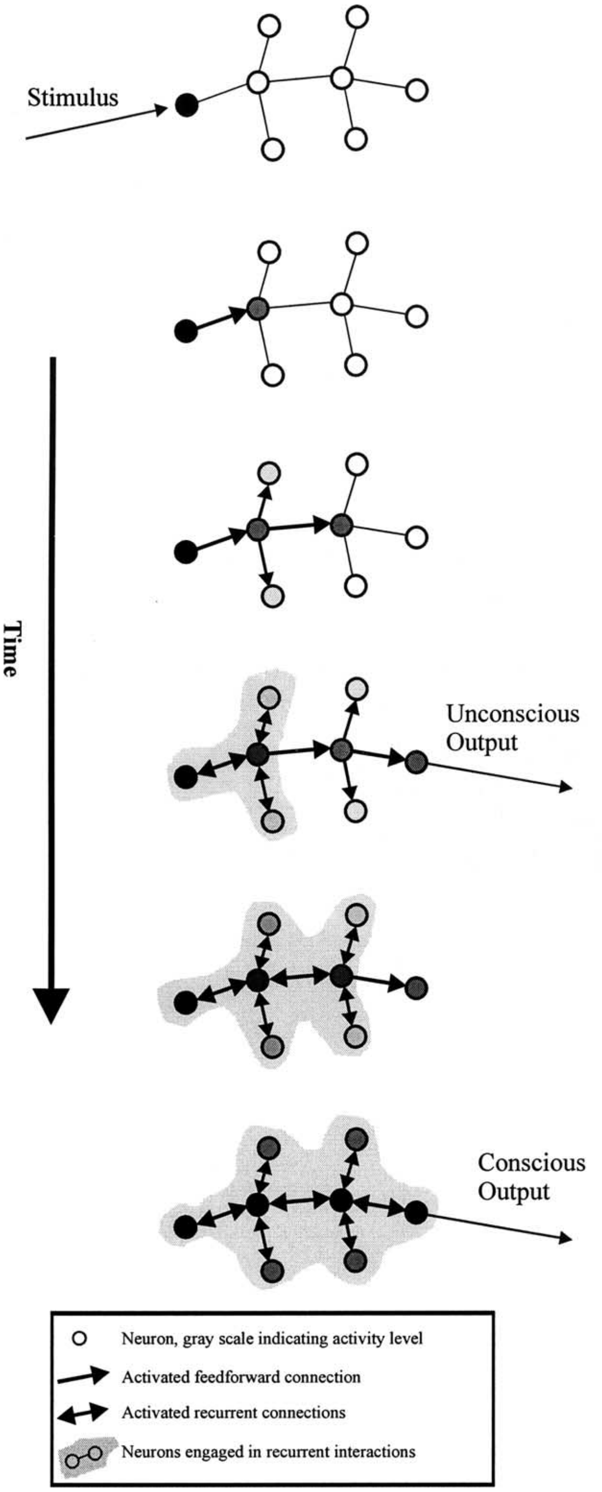

Fig. 4. Conscious visual experience requires recurrent processing. Feedforward processing rapidly activates (unidirectional arrows) successive levels of processing (from left to right), potentially leading to a reflex-like unconscious output or modification of behavior, based on basic (‘hard-wired’) categorizations and stimulus-response associations. Recurrent processing, mediated by horizontal or feedback connections (bi-directional arrows), lags behind this feedforward sweep (unless parallel feedforward sweeps exist, one being slower than the other, see Bullier, 2001). Recurrent processing mediates more complex stimulus-response associations, and is required for visual awareness, or for a conscious response.

thing), only weaker. Moreover, this activation is specific, e.g. distinguishes between invisible faces and houses (Moutoussis & Zeki, 2002). This shows that (feedforward) activation of cortical areas per se is insufficient for a conscious experience.

3. With transcranial magnetic stimulation (TMS) the ongoing activity in a particular brain region can be shortly disrupted. Applying TMS to early visual areas at a latency far beyond the FFS still renders stimuli invisible (Corthout, Uttl, Juan, Hallett, & Cowey, 2000). Also TMS over the motion selective area MT induces motion sensations, unless V1 activity is disrupted at a later moment in time (Pascual-Leone & Walsh, 2001). Since MT is higher in the visual hierarchy than V1, this implies that feedback from MT to V1 is necessary for motion awareness.
4. Feedforward activation of neurons can still be recorded in anesthetized animals, with RF tuning properties that hardly differ from those in the awake animal. Manifestations of recurrent processing, in particular those contextual modulations that express aspects of perceptual organization, are, however, reduced or fully suppressed under anesthesia (Lamme et al., 1998b).
5. Feedforward activation of neurons in V1 is not affected when stimuli are reported as not seen by animals engaged in a figure-ground detection task. A neural correlate of figure-ground segregation, probably mediated by recurrent interactions between V1 and extra-striate areas, and present when stimuli are seen, is, however, fully suppressed when stimuli are not seen (Super et al., 2001).

This has led me, and others, (Lamme, 2000, 2003) to conclude that visual processing mediated by the FFS, however, sophisticated, is not accompanied by conscious visual experience. Recurrent interactions are necessary for visual consciousness to arise (Fig. 4).

3.4. Consciousness × attentional selection: three stages of processing

We may now have a look (Fig. 5, see also Fig. 2f) at what happens when the proposed neural mechanism of visual consciousness (recurrent processing) interacts with the mechanism of attention (processing × memory). Suppose a visual scene being presented to the eyes. The feedforward sweep reaches V1 at a latency of about 40 ms. If multiple stimuli are presented, these are all represented at this stage. Next (60–80 ms), this information is fed forward to the extrastriate areas. At these intermediate levels, there is already some competition between multiple stimuli, in particular when they are close by. Not all stimuli can be processed in full by the receptive fields, that get larger and larger going upstream in the visual cortical hierarchy. This results in crowding phenomena. Attentional selection (in one way or another, see above), may resolve this competition (Desimone, 1998). In the end, only a few stimuli reach the highest levels, up to and including areas in executive space. This whole feedforward event evolves very rapidly (within ~120 ms) and is hypothesized to be fully unconscious. Feedforward activation alone may under certain circumstances result in a behavioral response (or modify ongoing behavior), but if it does, it will be a reflex-like action, that is fully unconsciously initiated (which is not to say that we may not become aware of it later).

Meanwhile, the early visual areas have started to engage in recurrent interactions, mediated by horizontal and feedback connections. By means of these recurrent interactions, visual features are related to each other, binding and segregation may occur, and perceptual organization evolves. This is what produces a conscious visual experience. Without these recurrent interactions there is no experience at all. Because at low levels there is relatively little competition between stimuli (unless they are close by), groups of recurrent interactions representing multiple stimuli are possible. This may occur for many items in a scene.

When these recurrent interactions grow more and more widespread, and eventually include areas in executive or mnemonic space (frontal, prefrontal, temporal cortex), the visual information is put into the context of the systems' current needs, goals, and full history. There is considerable competition, however, for interaction with these higher levels. Attentional selection during the feedforward sweep will already have predisposed some interactions over others. Alternatively, this selection may operate at the recurrent interactions themselves. In any case, only a limited number of recurrent groups can span the range from visual to more frontal (parietal, temporal) areas. Therefore, these more widespread recurrent interactions are limited to a few items in the scene. On the other hand, the recurrent process is less 'superficial' in the sense that stimuli are processed deeper; more behavioral and mnemonic context is added to the stimuli than in the case of more low level recurrent interactions, limited to the visual areas only.

We are thus able to discern at least three fundamentally different types of processing: (1) feedforward processing, influenced by attentional selection, but unconscious for reasons outlined above. (2) Recurrent processing of a restricted nature, limited to, say, visual areas. (3) Widespread recurrent processing of information that involves many regions of the brain and has passed the attentional bottleneck between sensory and executive areas. It is clear that widespread recurrent interactions should correspond to a stage of processing that we could call conscious. Here, a selected part of the information is embedded in mnemonic and behavioral context. Information thus processed is available for conscious access and can be reported about. Related theories refer to such a state as 'resonant' (Edelman, 1992; Grossberg, 1999), or as having reached 'global workspace' (Dehaene & Naccache, 2001). But what exactly is the nature of information that has achieved local (say only visual) recurrent embedding? It sits clearly between feedforward (unconscious) and globally recurrent (access conscious) processing, and for that reason alone it deserves a separate name. I here argue that locally recurrent processing is the neural correlate of phenomenal experience per se, or phenomenal awareness. That is a name, so the question is: what do I mean by that? This brings us back to where I left off at Section 2.5.

4. A case for phenomenal awareness

At the end of Section 2.5 I left with the question about what distinguishes unconscious from conscious visual processing. In the words of Section 3: what happens as processing evolves from feedforward to (at first locally) recurrent processing? What exactly happens when we cross the demarcation in Fig. 2d and f between unconscious and conscious.

Visual stimuli, or attributes of visual stimuli, that activate cortical neurons do not necessarily reach consciousness. Already mentioned examples are the high temporal and spatial frequency luminance patterns that we cannot see, yet that still excite V1 neurons (He & MacLeod, 2001). Other examples are patterns that are rendered invisible by some manipulation, such as crowding or masking (Enns & Di Lollo, 2000; Intriligator & Cavanagh, 2001). In the case of masking, it has been shown that these invisible stimuli evoke activity in visual areas as widespread as V1 (MacKnik & Livingstone, 1998), IT (Rolls & Tovee, 1994; Kovacs et al., 1995), the frontal eye fields (Thompson and Schall, 1999; 2000), and even the motor cortex (Dehaene et al., 1998). In the case of crowding, neural activation by invisible stimuli has been assumed on the basis of the fact that these stimuli are capable of producing adaptation effects (He, Cavanagh, & Intriligator, 1996). In all these situations, it can be successfully argued that feedforward processing has occurred for these stimuli or stimulus properties, yet that recurrent interactions are absent (Lamme, 2003; Lamme & Roelfsema, 2000; Lamme et al., 2000, 2002).

The conscious/unconscious distinction can be further refined: It is also very difficult, for example, to become aware of the physical wavelength of the light emitted by an object, instead of its perceived color, due to our color constancy mechanisms (Hurlbert, 1999). Another important insight comes from bi-stable or rivalrous patterns. These patterns can be viewed in either of two ways, and these alternate views last for durations that have a very characteristic distribution. Moreover, which of the two percepts are viewed is hardly under voluntary or attentional control (Leopold & Logothetis, 1999). This is a very strong case of a perceptual, rather than attentional, dichotomy that occurs while the physical retinal stimulation stays the same.

Finally, there are instances in which visual stimuli that are usually very well visible are not seen, even though the attentional and response systems are fully allocated.

For some reason, in those instances, sensory processing does not seem to complete to a stage that results in a percept of a particular object or scene property (Super et al., 2001, 2003). Where bistability or rivalry causes sensory processing to alternate between one percept and another, these situations cause sensory processing to alternate between a percept and no percept.

On the basis of these findings some line can be drawn between what we should call conscious and unconscious visual processing. Conscious visual stimuli have reached a level of processing beyond initial feature detection, where at least an initial coherent perceptual interpretation of the scene is achieved. Whether at this stage the binding problem (Driver, 1998), in all its diversity, has been solved is not clear at this point. The binding of some features of an object, such as its color and shape, may require attention, while other feature combinations are detected pre-attentively (Treisman, 1996; Driver, 1998; but see Di Lollo (2001)). So it may be that the conscious level, before attention has been allocated, consists of only tentatively (but uniquely) bound features, something that others have called proto-objects. There is a clear distinction, however, with unconscious stages, where individual features, even features that are never perceived, are represented.

During the feedforward sweep, information has been extracted, but this information has not yet interacted. Interaction between the distributed information requires recurrent interactions. Visual recurrent processing goes beyond initial feature detection and may be the neural correlate of binding or perceptual organization: features are tentatively bound, surfaces are defined, figure-ground relationships may be established. Others have called this stage 'mid-level vision' (Nakayama, He, & Shimojo, 1995) or the '2.5D sketch' (Marr, 1982). There is strong evidence that this stage is indeed a manifestation of recurrent interactions between early visual areas (Lamme & Spekreijse, 2000). Moreover, this stage is hardly susceptible to attentional bottlenecks or top-down control.

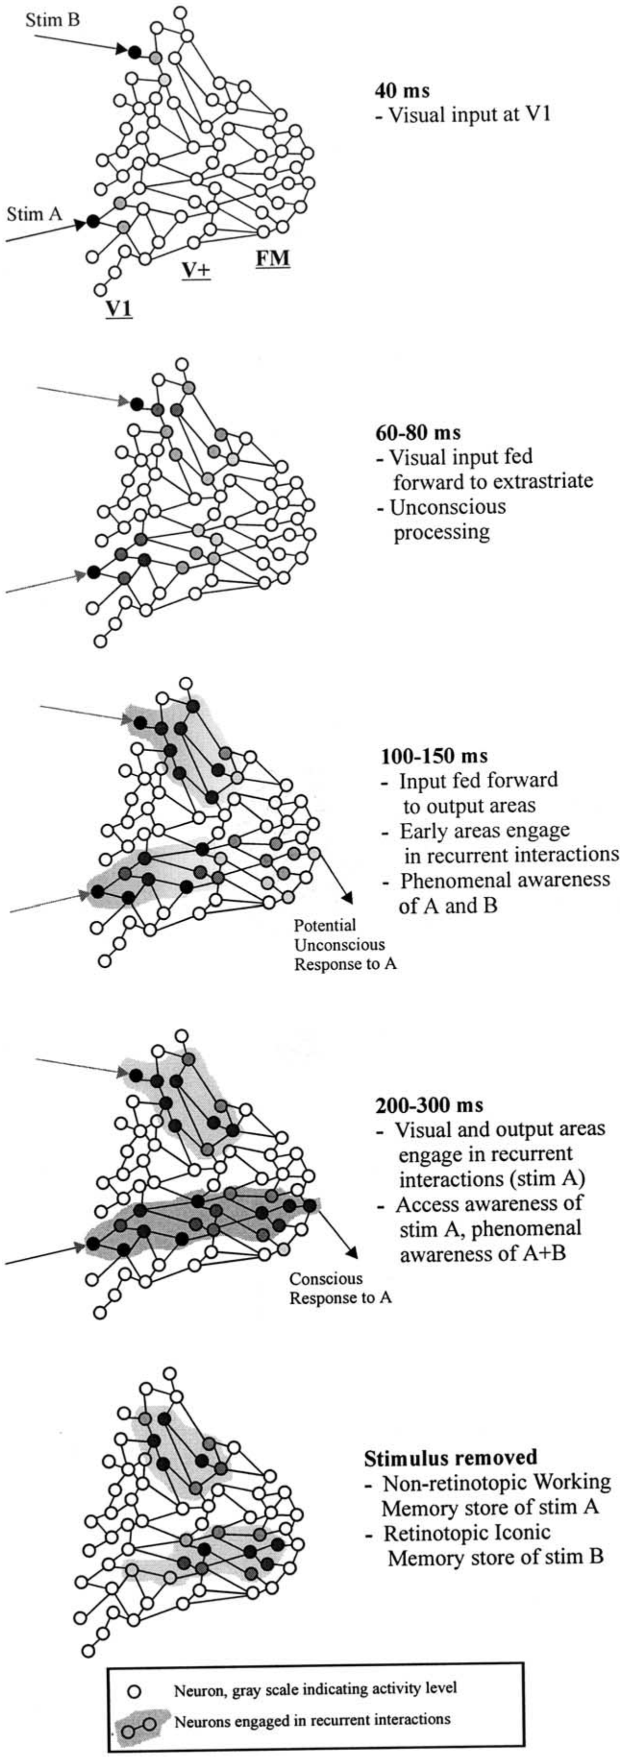

Fig. 5. Phenomenal versus access awareness. The interaction between recurrent processing (Fig. 2) and mechanisms of attentional selection (Fig. 1) is shown. As in Fig. 2, competition between the neural representations of multiple stimuli (stim A, B) may prevent the feedforward transfer from V1 to the executive areas (FM, frontal or motor regions) of all but a few stimuli (in this case A). At lower levels (V1, V +, extrastriate areas), however, simultaneous representations (of both A and B) may exist. Either way, feedforward activation (gray dots), both of selected (i.e. attended) and not selected inputs are unconscious, even though it may trigger or modify behavior. Meanwhile, neurons in activated regions may start to engage in recurrent interactions, which is accompanied by increased activity or synchronous firing (dots enclosed by gray shading). This produces *phenomenal* awareness of the visual inputs (and iconic memory after removal of the stimulus). Some of these recurrent interactions grow more widespread than others, and may even incorporate high level, executive, or planning regions, depending on attentional selection, in part already established during the feedforward sweep. Stimuli associated with these widespread interactions reach *access* awareness, and may be stored in non-retinotopic working memory after removal of the stimulus.

On the other hand, it is the first stage to which we have conscious access. Ingenious stereo display experiments have shown very elegantly that vision (more specifically, motion perception, object recognition and visual search) is based on this surface representation stage, and that it is the basis of our phenomenal experience (Nakayama et al., 1995). In my view, recurrent processing limited to the visual areas (Fig. 5) forms the neural basis of this stage, a stage where perceptual organization occurs, and that we should call phenomenally conscious (Fig. 2f).

To report about these percepts, however, the information has to become globally recurrent, has to reach access awareness. But that is not to say that before that we have no phenomenal experience of this information. We do, and that gives us the rich experience of vision we have. This experience is not an illusion, as change blindness (Simons & Levin, 1997) or inattentional blindness (Mack & Rock, 1998) experiments might suggest (O'Regan & Noe, 2001). Those experiments reveal the attentional bottleneck between experience and report (Landman et al., 2003), or between experience and the storage of experiences (Wolfe, 1999), not the limitations of experience itself. The distinction between phenomenal and access awareness that has been made by Ned Block (1996) on philosophical and theoretical grounds is very related to this. Also in the domain of memory we find a similar distinction. Working memory, the limited capacity, yet stable storage of information (Cowan, 1994), has clear similarities to access awareness, and may even share neural mechanisms. Iconic memory, the large capacity, yet fleeting form of memory we have of a scene in its entirety (Coltheart, 1980), may be linked to phenomenal awareness. Both are forms of memory, however, and in that sense are different from conscious visual experience, which is only present when stimuli are there.

5. Conclusions

From the cognitive neuroscience perspective a clear distinction can be made between visual attention and visual consciousness. Attentional selection is how sensorimotor processing is modified by the current state of the neural network, shaped by genetic factors, experience, and recent events (memory). Phenomenal experience has a different origin, which is the recurrent interaction between groups of neurons. Depending on the extent to which recurrent interactions between visual areas incorporate interactions with action or memory related areas, awareness evolves from phenomenal to access awareness. Whether this occurs depends on attentional selection mechanisms, via influences on both the feedforward sweep and recurrent interactions. Other mechanisms, however, determine whether neurons will engage in recurrent interactions at all, and thus whether processing will go from an unconscious to a conscious state.

The hypothesis forms a core understanding of the different forms of visual consciousness, in the same spirit as we have core understandings of sensori-motor transformations, memory and attentional selection. Again, many things still need to be worked out, the most important being of course to explain why recurrent interactions are necessary for phenomenal experience to arise, and how we go from such a neural process to the phenomenon of mental experience. In that sense, the 'hard problem' remains as hard as it was. With this theory, however, we may have a better sense of what to look for.

References

Albright, T. D., & Stoner, G. R. (2002). Contextual influences on visual processing. *Annual Review in Neuroscience*, *25*, 339–379.

Becker, M. W., et al. (2000). The role of iconic memory in change detection tasks. *Perception*, *29*, 273–286.

Block, N. (1996). How can we find the neural correlate of consciousness. *Trends in Neuroscience*, *19*, 456–459.

Bullier, J. (2001). Integrated model of visual processing. *Brain Research Reviews*, *36*, 96–107.

Colby, C. L., & Duhamel, J. R. (1996). Spatial representations for action in parietal cortex. *Brain Research Cognition Brain Research*, *5*, 105–115.

Coltheart, M. (1980). Iconic memory and visible persistence. *Perception and Psychophysics*, *27*, 183–228.

Corthout, E., Uttl, B., Juan, C. H., Hallett, M., & Cowey, A. (2000). Suppression of vision by transcranial magnetic stimulation: a third mechanism. *Neuroreport*, *11*, 2345–2349.

Cowan, N. (2001). The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behavioral Brain Science*, *24*, 87–185.

Crick, F., & Koch, C. (1998). Consciousness and neuroscience. *Cerebral Cortex*, *8*, 97–107.

Cumming, B. G., & Parker, A. J. (1997). Responses of primary visual cortical neurons to binocular disparity without depth perception. *Nature*, *389*, 280–283.

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition*, *79*, 1–37.

Dehaene, S., et al. (1998). Imaging unconscious semantic priming. *Nature*, *395*, 597–600.

Desimone, R. (1996). Neural mechanisms for visual memory and their role in attention. *Proceedings of the National Academy of Sciences USA*, *93*, 13494–13499.

Desimone, R. (1998). Visual attention mediated by biased competition in extrastriate visual cortex. *Philosophical Transactions of the Royal Society of London, Series B*, *353*, 1245–1255.

Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review in Neuroscience*, *18*, 193–222.

Di Lollo, V., et al. (2001). The preattentive emperor has no clothes: a dynamic redressing. *Journal of Experimental Psychology General*, *130*, 479–492.

Driver, J. (2001). A selective review of selective attention research from the past century. *British Journal of Psychology*, *92*, 53–78.

Driver, J., & Frackowiak, R. S. (2001). Neurobiological measures of human selective attention. *Neuropsychologia*, *39*, 1257–1262.

Dudai, Y. (2002). Molecular bases of long-term memories: a question of persistence. *Current Opinion in Neurobiology*, *12*, 211–216.

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

Egeth, H. E., & Yantis, S. (1997). Visual attention: control, representation, and time course. *Annual Review in Psychology*, *48*, 269–297.

Engel, A. K., et al. (2001). Dynamic predictions: oscillations and synchrony in top-down processing. *Natural Review in Neuroscience*, *2*, 704–716.

Enns, J. T., & Di Lollo, V. (2000). What's new in visual masking? *Trends in Cognitive Science*, *4*, 345–352.

Fries, P., et al. (2001). Modulation of oscillatory neuronal synchronization by selective visual attention. *Science*, *291*, 1560–1563.

Goldman-Rakic, P. S. (1996). Regional and cellular fractionation of working memory. *Proceedings of the National Academy of Sciences USA*, *93*, 13473–13480.

Grossberg, S. (1999). The link between brain learning, attention and consciousness. *Consciousness and Cognition*, *8*, 1–44.

He, S., Cavanagh, P., & Intriligator, J. (1996). Attentional resolution and the locus of visual awareness. *Nature*, *383*, 334–337.

He, S., & MacLeod, D. I. (2001). Orientation-selective adaptation and tilt after-effect from invisible patterns. *Nature*, *411*, 473–476.

Hurlbert, A. (1999). Colour vision: Is colour constancy real? *Current Biology*, *9*, R558–R561.

Intriligator, J., & Cavanagh, P. (2001). The spatial resolution of visual attention. *Cognitive Psychology*, *43*, 171–216.

Kovacs, G., Vogels, R., & Orban, G. A. (1995). Cortical correlate of pattern backward masking. *Proceedings of the National Academy of Sciences USA*, *92*, 5587–5591.

Lamme, V. A. F. (2000). Neural mechanisms of visual awareness: a linking proposition. *Brain and Mind*, *1*, 385–406.

Lamme, V. A. F. (2003). Why visual attention and awareness are different. *Trends in Cognitive Sciences*, *7*, 12–18.

Lamme, V. A. F., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences*, *23*, 571–579.

Lamme, V. A. F., & Spekreijse, H. (2000). Modulations of primary visual cortex activity representing attentive and conscious scene perception. *Frontiers in Bioscience*, *5*, D232–D243.

Lamme, V. A. F., et al. (1998a). Feedforward, horizontal, and feedback processing in the visual cortex. *Current Opinion in Neurobiology*, *8*, 529–535.

Lamme, V. A. F., et al. (1998b). Figure-ground activity in primary visual cortex is suppressed by anaesthesia. *Proceedings of the National Academy of Sciences USA*, *95*, 3263–3268.

Lamme, V. A. F., et al. (2000). The role of primary visual cortex (V1) in visual awareness. *Vision Research*, *40*, 1507–1521.

Lamme, V. A. F., et al. (2002). Masking interrupts figure-ground signals in V1. *Journal of Cognitive Neuroscience*, *14*, 1–10.

Landman, R., et al. (2003). Large capacity storage of integrated objects before change blindness. *Vision Research*, *43*, 149–164.

Leopold, D. A., & Logothetis, N. K. (1999). Multistable phenomena: changing views in perception. *Trends in Cognitive Sciences*, *3*, 254–264.

Levy, R., & Goldman-Rakic, P. S. (2000). Segregation of working memory functions within the dorsolateral prefrontal cortex. *Experimental Brain Research*, *133*, 23–32.

Mack, A., & Rock, I. (1998). *Inattentional blindness*. Cambridge, MA: MIT Press.

Marr, D. (1982). *Vision*. New York: W.H. Freeman.

Moutoussis, K., & Zeki, S. (2002). The relationship between cortical activation and perception investigated with invisible stimuli. *Proceedings of the National Academy of Sciences USA*, *99*, 9527–9532.

Nakayama, K., He, Z. J., & Shimojo, S. (1995). Visual surface representation: a critical link between lower-level and higher level vision. In S. M. Kosslyn, & D. N. Osherson (Eds.), *Vision. In invitation to cognitive science* (pp. 1–70). Cambridge, MA: MIT Press.

O'Regan, J. K., & Noe, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral Brain Science*, *24*, 939–1031.

Oram, M. W., & Perrett, D. I. (1992). The time course of neural responses discriminating between different views of the head and face. *Journal of Neurophysiology*, *68*, 70–84.

Pascual-Leone, A., & Walsh, V. (2001). Fast backprojections from the motion to the primary visual area necessary for visual awareness. *Science*, *292*, 510–512.

Posner, M. I. (1994). Attention: the mechanisms of consciousness. *Proceedings of the National Academy of Sciences USA*, *91*, 7398–7403.

Rensink, R. A. (2000). Seeing, sensing, and scrutinizing. *Vision Research*, *40*, 1469–1487.

Rensink, R. A. (2002). Change detection. *Annual Review in Psychology*, *53*, 245–277.

Rolls, E. T., & Tovee, M. J. (1994). Processing speed in the cerebral cortex and the neurophysiology of visual masking. *Proceedings of the Royal Society of London series B, Biological Science*, *257*, 9–15.

Searle, J. R. (1998). How to study consciousness scientifically. *Philosphical Transactions of Royal Society of London B Biological Sciences*, *353*, 1935–1942.

Simons, D. J. (2000a). Attentional capture and inattentional blindness. *Trends in Cognitive Science*, *4*, 147–155.

Simons, D. J. (2000b). Current approaches to change blindness. *Visual Cognition*, *7*, 1–15.

Simons, D. J., & Levin, D. T. (1997). Change blindness. *Trends in Cognitive Science*, *1*, 261–267.

Super, H., et al. (2001). Two distinct modes of sensory processing observed in monkey primary visual cortex (V1). *Natural Neuroscience*, *4*, 304–310.

Super, H., van der Togt, C., Spekreijse, H., & Lamme, V. A. F. (2003). Internal state of monkey primary visual cortex (V1) predicts figure-ground perception. *Journal of Neuroscience*, *23*, 3407–3414.

Thompson, K. G., & Schall, J. D. (1999). The detection of visual signals by macaque frontal eye field during masking. *Nature Neuroscience*, *2*, 283–288.

Treisman, A. (1996). The binding problem. *Current Opinion in Neurobiology*, *6*, 171–178.

VanRullen, R., & Thorpe, S. J. (2001). Is it a bird? Is it a plane? Ultra-rapid visual categorisation of natural and artifactual objects. *Perception*, *30*, 655–668.

Wickens, T. D. (2002). *Elementary signal detection theory*. New York: Oxford University Press.

Wolfe, J. M. (1999). Inattentional amnesia. In V. Coltheart (Ed.), *Fleeting memories*. Cambridge: MIT Press.

Zeki, S., & Marini, L. (1998). Three cortical stages of colour processing in the human brain. *Brain*, *121*, 669–685.