John Benjamins Publishing Company

Neuronal reflections and subjective awareness

Rafael Malach

Department of Neurobiology, Weizmann Institute of Science

For a neuroscientist working under the assumption of a complete correspondence between mind and brain, conscious awareness poses a profound mystery. It is a unique phenomenon in which a state of a physical system – i.e. a distributed pattern of neuronal activity is inexplicably transformed into an absolutely private and internal mental experience. Thus, the relevant point of view that needs to be taken when considering neuronal mechanisms underlying phenomenal experience must also be intrinsic – that of the brain's neurons themselves. In contrast, the perspective taken by most neuroscience research examines neuronal activity from an external observer perspective. Here I will conjecture that the intrinsic perspective is implemented in the brain through rapid and recurrent neuronal activity – local neuronal "ignitions". In this dynamics, the information about the state of a local neuronal assembly is distributed back to the neurons that form the assembly through recurrent activations. A conscious percept emerges when, through these neuronal reflections, the ambiguity inherent in the activity of isolated neurons is converted into a unique and meaningful assembly state. The rapid distribution of assembly information necessitates high firing rates, sustained activity, and dense local connectivity. All these conditions fit nicely with recent experimental findings. The fourth consequence – the founding of conscious awareness on local reverberatory activity – is still highly controversial – and should be viewed, at this stage, as a prediction of the local "ignition" hypothesis.

1.   Introduction – the optimistic outlook

It is a self-evident fact that our subjective, phenomenal state, e.g. the image we have when marveling at the starry night, is the most direct and accessible of all occurrences we can conceive of. Thus, one would expect that this universe of sensations will be high on the priority list of neuroscientists to explore and explain. Yet, the contrary is true – the scientific exploration of phenomenal experience has been extremely hesitant and plagued by controversies. Thus, it is still debated whether subjective experience is at all amenable to proper scientific research. The reason is of course the very unique nature of subjective experience – which drastically differs from all other natural phenomena. Particularly problematic to

scientific research is the private nature of subjective experience which is inaccessible to direct measurements or objective observation.

Furthermore, unlike all other physical phenomena – the link of subjective experience to the physical world is deeply unintuitive and in fact incomprehensible. It is an enigmatic process by which a physical system – i.e. a collection of nerve cells in an individual's brain – is somehow capable of "becoming aware" – i.e. transforming a distributed pattern of neuronal activity into an absolutely private and unitary phenomenal experience. Note that this statement makes two implicit assumptions, which should be stated upfront. First, it is assumed that phenomenal experience is indeed linked to the brain – the philosophical stance I will hold is that of monism – as formulated already by Spinoza' – i.e. that subjective experience and the physical state of the brain are two different perspectives of one and the same phenomenon. The second assumption is that there is a one to one mapping between contents of phenomenal experiences and physical states of the brain. Discussing alternatives to these basic assumptions will be beyond the scope of this chapter.

It has been extensively debated whether the unique nature of phenomenal experience may actually preclude the construction of a scientific theory of conscious awareness – which is after all the goal of any scientific endeavor. Here I will adopt a more optimistic outlook and will argue that substantial progress in the scientific search for the neuronal basis of subjective experience is feasible. However it necessitates a major shift in our point of view when considering models of phenomenal awareness. If we accept that subjective experience entails some internal link to the state of the brain's neurons – then as scientists trying to understand this process, we must forego our natural perspective as experimentalists – that of external observers collecting information about the external reality, and instead adopt an "intrinsic" perspective. Such perspective should examine what information about the state of the brain is actually available to its constituent neurons on their own – unaided by the knowledge we experimentalists gain in the course of our observations.

As an example for subjective phenomenal experience I will use conscious visual perception – since my research has mainly focused on this type of experience. But it should be possible to generalize the considerations derived from this research to the entire set of phenomenal subjective experiences.

2.   The creative nature of visual perception

Starting with Helmholz and the Gestalt Psychologists – our understanding of the behavioral aspects of visual perception has shown a tremendous progress leading to the establishment of visual psychophysics as a rigorous and rich field of active

research. Here I will focus only on few general principles that this research has led to and which are directly relevant to the argument I am trying to develop. One such fundamental principle is the understanding that the content of a conscious visual percept (which I will operationally define as any visual experience that a healthy individual can both introspectively and objectively report about) can often go far beyond the immediate optical or retinal information provided by the eyes. Two classic examples that illustrate this point are depicted in Figure 1. In the phenomenon of illusory contours (e.g. the Kanizsa figure, panel A) and in the case of "impossible figures" – panel B. In both cases the phenomenal percept actually *contradicts* the physical objective information. The physical world and the retinal image consist of isolated wedges in the case of the Kanizsa triangle, yet the phenomenal percept contradicts this objective information and generates instead a subjective experience of a triangle which does not exist in reality. While in the case of the illusory triangle one may still invoke prior experience or memory with such triangles as the basis of the "illusion" – in the case of the impossible figure even this can not be true – since such objects can not possibly exist in the real world – and hence, surely were not experienced by the observer prior to seeing them.

Illusory triangle

Impossible triangle
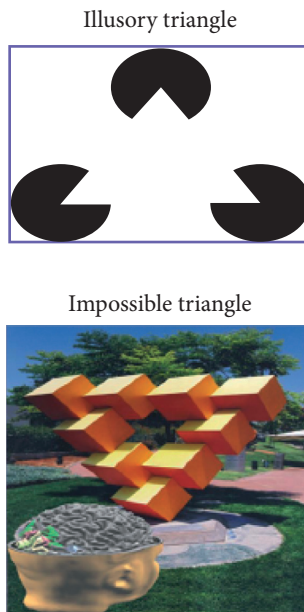
Figure 1. Two visual illusions illustrating the discrepancy between visual perception and objective reality. The illusory triangle is visible despite the physical absence of edges. The "impossible" triangle is easily perceived as a triangle despite the fact that such objects can not possibly exist in the physical world

These are of course just two examples out of a large set of striking illustrations – all pointing to an unavoidable conclusion – the visual percept can not be viewed as some kind of transformation, processing or "representation" of the physical visual information. A straightforward consequence of this is that the information that underlies the creation of a phenomenal visual image must be available internally – within the brain itself. The most striking illustration of this fact is in the spontaneous emergence of vivid visual images during dreaming as well as following electrical stimulation in patients, in visual hallucinations or upon trans-magnetic stimulation (TMS). In all these cases, a vivid visual percept is generated in the complete absence of any optical stimulation or an external visual target. This is not meant to imply that the external world is irrelevant to perception – to the contrary – clearly a large part of the internal information underlying the phenomenal image reflects regularities and adaptive values of the external physical world. The point is that during a moment of a perceptual awareness, when we see a vivid visual image, the information pertaining to this image is derived mainly from the brain itself. Thus, it is in the brain, not in the outside world, that we should search for the source of information allowing the moment by moment creation of visual percepts.

3. A colossal library

A phenomenal visual percept can not therefore be viewed as some kind of transformation of optical information derived from the external world, but is rather derived from information stored internally in the visual system itself.

The significance of this conclusion becomes evident when we consider the sheer number of different visual percepts that a single human visual system is capable of generating. Taking just the case of faces, a person is quite likely capable of differentiating between most faces belonging to one's own race – likely millions of different faces. Considering the fact that each of these individual faces can appear in different colors, sizes, contrasts, expressions, ages etc. – it becomes clear that our capacity to generate differentiable visual contents runs in the many trillions of different images.

It is important to clarify that we are not dealing here with invariant visual recognition or memory – where the capacity is likely far more limited. The challenge for a theory of phenomenal visual experience concerns the staggering amount of potential visual images – i.e. we need to explain how brain states can be constructed so that such an enormous set of subjective percepts can be differentiated. It is obvious but needs stressing, that when it comes to estimating our capacity to generate conscious visual percepts – any slight consciously perceived

changes – such as color changes of the same person – must give rise to different brain states. With this in mind, it is clear that the image generating capacity of the brain is truly colossal. So how are neurons capable of generating such a massive number of differentiable visual percepts?

4.   Building templates – hierarchically

In trying to come up with a solution, I will focus on the primate cortex – since it is in primates that we assume that the link between behavioral report and subjective experience is most reliable. Two points should be noted before we embark on this analysis. First, the fact that it may be more convenient to study primate vision does not imply that other animals have a reduced level of phenomenal visual experience. Second, even in humans, the ultimate inaccessibility of external observation to the subjective content of experience should be acknowledged – i.e. we can never be completely certain whether the report we receive truly reflects the phenomenal experience. Nevertheless, as I will try to illustrate below, some conceptual advances can be made with regards the relationship between subjective states and brain states.

Examining how the processing of visual information begins in the primate cortex, the first principle that becomes apparent is the hierarchical nature of the processing. Thus, converging evidence supports the notion that as we move anteriorly in the visual cortex, along the sequence of visual areas, the tuning of cortical neurons shifts from local, simple elements – such as lines, corners and contrast boundaries to more abstract and complex responses. The characteristic of neuronal tuning at the top of the visual hierarchy is still debated, but a growing body of work points to a strong holistic and configural element in the neuronal tuning. Thus, both human brain imaging and monkey single unit recordings reveal selectivity to various visual categories. Furthermore, direct demonstrations of holistic and configural properties have been consistently documented (e.g. Grill-Spector & Malach 2004). All these results converge in supporting the notion of a holistic, template – selectivity of neurons at the top of the cortical hierarchy.

The importance of these observations is that they offer a straightforward (although far from simple) feed-forward mechanism that could potentially solve the classic "binding" problem. In such a formulation, each neuron at the top of the hierarchy binds the local visual elements through converging feed-forward inputs from lower level visual areas. Thus, simple visual elements represented by the activity in primary visual neurons are ultimately "bound" into holistic templates corresponding to the unified perceptual images at the top of the visual hierarchy.

This is the classic sparse representation – popularly coined the "grand-mother" neuron hypothesis (Barlow 1972) – that keeps emerging under various modeling attempts.

5. A combinatorial explosion

While feed-forward hierarchical convergence could potentially solve the binding problem – it does not address the other fundamental problem raised earlier – the capacity of the visual system to generate trillions of distinguishable perceptual images. Note that even if we allocate a single "grand-mother" neuron to each perceptual image – there will not be enough neurons in the brain to account for the vast collection of all distinguishable conscious percepts. The solution, as indeed has been proposed by a number of distributed models, is that each neuron will be sensitive to a large number of different visual templates. We have previously proposed the metaphor of a "totem-pole" neuron to illustrate this notion (Gilaie-Dotan & Malach 2007). Basically the idea is that a neuron's tuning can be envisioned as an or-function of many, likely millions, of different distinguishable holistic image templates – so that whenever we perceive a specific image, all the neurons that include this specific template in their tuning profile will fire.

The rules that organize the library of visual templates within each individual "totem-pole" neuron are not fully established. The discovery of face neurons as well as other object category neurons, suggests that the totem-pole libraries are organized according to a categorical principle. Furthermore, brain imaging and direct recordings from human cortex suggest that perceptual similarity may be another underlying principle – e.g. similarly looking exemplars within a category are more likely to be members of the same "totem-pole" library – i.e. co-activate a more similar set of neurons as compared to non-similar exemplars.

6. The ambiguity inherent in isolated responses

Regardless of the precise nature of the totem-pole tuning curves, the important principle for our consideration here is the fact that given the limited number of visual neurons – and the vast space of potential conscious percepts – each visual neuron must be tuned to a very large set of different visual percepts.

The problem with such broadly tuned neurons is that they are individually ambiguous – i.e. when an individual neuron is firing it is not possible to know which specific perceptual image this firing refers to. The solution to this

ambiguity, as indeed has been proposed by a number of models is in assembly representation. Thus, to take the totem-pole metaphor – if all neurons that contain a single common image, x, in their library are active together – this co-activation will unambiguously "point" to x as the perceived image. In contrast, those images which some of the neurons in the assembly are not sensitive to – will be ruled out.

Importantly, the library of each totem-pole neuron contains a subset of all the possible visual images, and furthermore, different totem-pole neurons have overlapping, yet different libraries. Such an arrangement generates a huge combinatorial space of possible images. Thus, if we consider the case of face representations, we have estimated indirectly that at least a million neurons are linked to the perception of each individual face (Levy, Hasson & Malach 2004). The number of different combinations that can be generated in assemblies consisting of a million neurons is truly colossal – and is more than sufficient to account for all possible perceptual images.

However, the principle of assembly coding of holistic "totem-poles", which seems to solve both the problem of binding and that of the need for a huge storage of perceptual images, leads to the fundamental, often overlooked, problem of readout. As I have argued earlier, any model of conscious perception must consider the intrinsic perspective of the brain itself and not that of the experimenter collecting the data. The question that must be resolved then is that: what system within the brain itself, actually "knows" that a specific group of neurons – e.g. the assembly representing image x – has been activated?

7.   Neuronal reflections

One can of course envision numerous mechanisms by which the activity of neuronal groups can be detected in the brain – however there is one unique process of information exchange that implements the self-readout concept in the most literal manner. I would like to put forward the conjecture that such process is the core dynamics underlying all conscious awareness states.

More specifically, I proposed that a visual percept emerges when the state of an activated assembly is made available to each of the member neurons that take part in that assembly. In other words, there is a special state in which each neuron in an assembly receives, directly or indirectly, a pattern of synaptic activations that reflect *a single and unique* activation state of all neurons in the assembly. Critically, the same single activation state is "felt" by all neurons simultaneously. This may immediately raise the worry of infinite regress – since the neurons that are affected by the state of the assembly are also the ones that constitute this state. However, as

has been demonstrated in many network models, such dynamics can settle, at least temporarily, into a stable consensus, in which all neurons are both affected and yet at the same time activate other members in the assembly. Note on the other hand, that such fully "democratic" distribution of information about a network state to all its members can be accomplished only in a recurrent network. By contrast, in feed-forward chains, the state of the system is not reflected in the synaptic inputs to all its member neurons.

To illustrate this conjecture in the context of visual perception – as was discussed above, the information inherent in the activity of isolated neurons is ambiguous and hence meaningless. However this ambiguity can be resolved by forming an integrated assembly. Here it is proposed that the means to achieve such critical integration is a unique dynamics in which the pattern of activity in the assembly is distributed to all its member neurons through reverberatory activity. Furthermore, when such recurrent dynamics reaches a momentary stable consensus – a perceptual image is created in the mind of the observer. We have recently proposed that such mutual activation among neighboring neurons can be metaphorically envisioned as a process of "ignition" (Fisch et al. 2009).

8.   Experimental consequences

Although the ignition conjecture requires detailed modeling which is beyond the scope of the present chapter – even considering its more qualitative aspects can provide well-defined constraints and predictions. An important consequence of the hypothesis concerns the time it takes to distribute the information that an assembly has been activated to all its member neurons. Note that under the plausible assumption that neuronal signaling is encoded in inter-spike time intervals (instantaneous firing rates) – the speed by which neurons can transmit information about the activation of an assembly is critically dependent on the firing rate. Thus, a system in which the transition from inactive to active state consists of a change in firing from 1 to 5 Hz – will take a minimum of 200 msec for one signal iteration – a very long time when considering the fast response time typical of perceptual processes (Hung, Kreiman, Poggio & DiCarlo 2005; Thorpe 2005). Thus, high firing rate is critical for such assembly coding if it needs to comply with realistic speeds of perceptual responses.

A second consequence relates to connectivity – if information among neurons must be distributed rapidly – there should be dense anatomical connectivity between them to allow such rapid distribution. Furthermore, the critical information exchange should occur among those neurons that must

share information – e.g. have common "libraries". Such neurons tend to cluster together anatomically – so that the connectivity must be mainly local, among close anatomical neighbors.

A third important consequence concerns stability – note that the content of the perceptual state is equivalent to the specific pattern of active neurons. To achieve a non-ambiguous percept, this pattern must remain stable for the duration it takes the neurons in the assembly to reach a consensus. This means that high firing rates must be sustained, and that a pair of spikes is not sufficient to signal the appearance of an assembly.

These are theoretical consequences of the ignition conjecture. The critical question is to what extent experimental findings confirm to these consequences. Below I will try to show that the ignition conjecture can account for a number of seemingly unrelated experimental observations from human and animal research.

9. What is so unique about the cerebral cortex?

The first experimental fact that needs to be taken into consideration is that the part of the brain which appears most closely associated with subjective awareness is the cerebral cortex. Thus both cortical lesions and electrical stimulation, and more recently TMS, have provided ample evidence to the critical link between this special organ and conscious awareness. But what makes the cerebral cortex so unique in its association with awareness phenomena? Below I will review the three most prominent aspects of cortical functional organization that emerge from recent neuroscience data.

Connectivity: Cortical neurons have dense local interconnections

A potential clue into the necessary conditions for awareness may be highlighted by examining what neuroanatomical features distinguish the cerebral cortex from other brain structures. While the feature that is classically emphasized is the laminated structure of the cortex, more recent neuroanatomical studies have also emphasized the dense recurrent connectivity (Douglas & Martin 2007). A powerful substrate for intense reverberatory activity, which has not received appropriate attention, is the dense "halo" of local connections which is particularly prominent in high order association cortex (See Figure 2, and (Amir, Harel & Malach 1993)) and contains numerous functional synapses. Such halos provide a dense substrate for reverberatory interactions at local scales – i.e. among neighboring neurons.
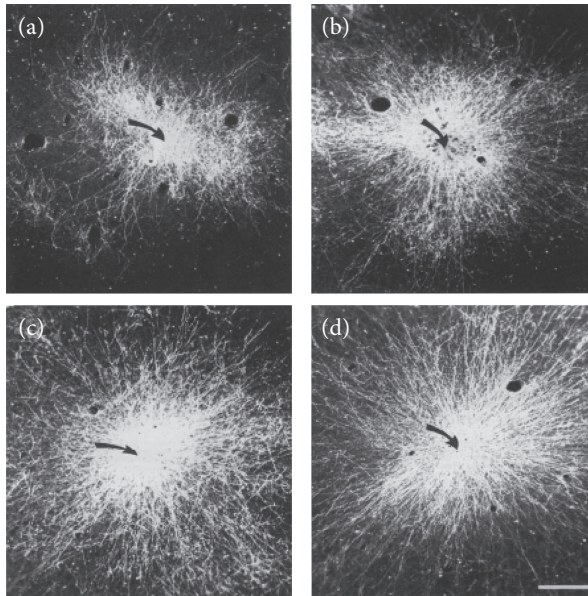
Figure 2. Top views of dense "halos" of lateral connections in primate visual cortex. Lateral connections can be seen emanating from neurons at the center of the tracer injection sites (arrows). Note the dense connectivity which increases in size from area V1 (a) through V2 (b) V4 (c) to 7a (d). The density of reciprocal interconnections can be appreciated when considering that such connectivity streams from every cortical site. Bar = 100 μm. (Taken from Amir et al. 1993)

The crucial question is whether there is any evidence that ties such lateral connections to conscious awareness? Perhaps the most dramatic demonstration of the critical role of intrinsic cortical connectivity comes from split brain patients. In these individuals, the surgical interruption of connections between neurons located in the two hemispheres literally splits (or doubles) the subjective awareness into two separate, autonomous entities. This truly amazing manipulation of the awareness state has been documented in the classic studies of Gazzaniga and others, and provides a strong indication that the unity of consciousness depends on the availability of intra-cortical neuroanatomical connections. Furthermore, although it may seem that cross hemispheric connections represent long-distance neuronal interactions – this is not the case. Both their laminar distribution as well as their anatomical locations (e.g. the vertical meridian representation in retinotopic cortex) indicate that the inter-hemispheric callosal connections can be largely viewed as an extension of the local, intrinsic connectivity, that has been "stretched out" as a consequence of the of dividing cortical areas into two anatomically separate hemispheres.

High Activity: Subjective awareness is associated with high neuronal firing rates.

Recent single unit and brain imaging studies have revealed a close link between perceptual awareness and intense neuronal activity in high order occipito-temporal cortical areas. Most direct are those studies involving bi-stable stimuli, in which the physical stimulus is held constant while subjective percepts alternate between two states. Thus, In the Rubin vase-face illusion (Hasson, Hendler, Ben Bashat & Malach 2001) a clear correlation has been demonstrated between the emergence of face percepts and the level of BOLD activity in occipito-temporal cortex. Both single units (Sheinberg, Leopold & Logothetis 1995) and brain imaging (Tong, Meng & Blake 2006) have revealed a link between high neuronal activity and perceptual state during binocular rivalry. Finally, using a backward masking paradigm, brain imaging (Grill-Spector, Kushnir, Hendler & Malach 2000) and intra-cranial ECoG recordings (Fisch et al. 2009) reveal that high neuronal activity is tightly linked to crossing the consciousness threshold. These are just a few examples out of a long list of studies that lead to the conclusion that at least in high order visual cortex high levels of neuronal activity are tightly linked to the emergence of perceptual awareness.

Time: Subjective awareness requires a minimal duration (100–300 msec) of neuronal activity.
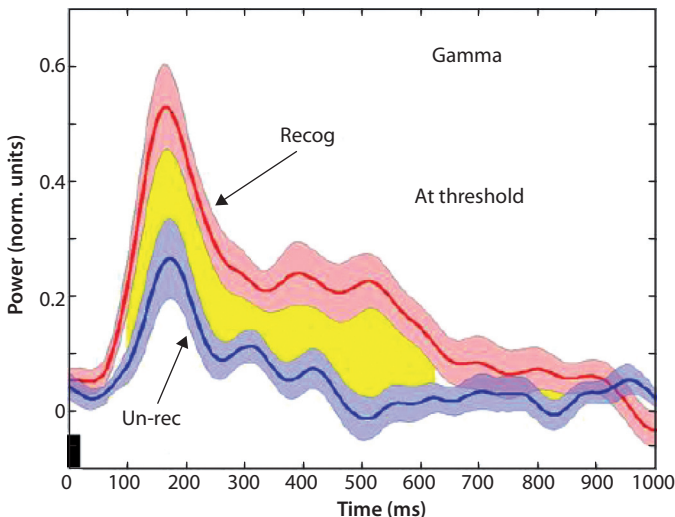
Figure 3. ECOG recording in high order visual cortex during visible and non visible trials in a backward masking experiment. Note the high and sustained neuronal activity during the visible (recog) state compared to the short and low amplitude response during the non-visible (Un-rec) trials. (Modified from Fisch et al. 2009)

Several lines of evidence point to the fact that subjective awareness can emerge only when neurons are active for a sufficient duration of time. Evidence for a minimal duration of neuronal activity comes from early studies of Libet, which demonstrated that, in patients undergoing surface stimulation of their primary somato-sensory cortex, sensory awareness emerged only after a rather prolonged (500 msec) period of cortical activation.

These results are also compatible with numerous visual masking experiments, in which the image is presented briefly, followed by a masking stimulus. These studies demonstrate that unlike non-visible stimuli, which lead to short lived neuronal activity, the neuronal responses associated with subjective awareness are typically sustained for 200–300 msec (Fisch et al. 2009). These results point to a minimal time that the neuronal signals need to "percolate" in the system to allow the perceptual event to be generated.

To summarize, converging experimental data point to high activity, dense connectivity and relatively long durations as critical for the emergence of perceptual awareness. These are precisely the same constraints that have been implicated as also critical for the formation of a consensus in neuronal assemblies proposed earlier. However, a fourth constraint – which is the most counter-intuitive and controversial concerns the locality of such assemblies.

10.   Is consciousness a local or global phenomenon?

The crux of the assembly ignition hypothesis is that it allows for the possibility that subjective awareness could emerge locally – even within a small cluster of interacting neurons. Thus, as far as subjective awareness proper is concerned, the hypothesis does not require that the consensus about the assembly state will include large scale networks outside of the active assembly proper. The notion of a local assembly consensus as sufficient for perceptual awareness is of course quite counter intuitive – and in its purest form, can not in fact be verified, since the only means for the external world to be informed about the emergence of subjective awareness is through some kind of motor output – which by necessity requires that information, e.g. from the visual cortex, will eventually be transferred to some kind of behavioral output mechanism (Malach 2007).

The question, however, can be presented in a less extreme form: thus, we can consider whether the hypothesized cortical "ignition" underlying subjective awareness is a local or global phenomenon (Dehaene & Changeux 2011). To go back to the example of visual awareness – the question can be formulated more concretely as to whether the emergence of a conscious visual image entails an ignited assembly that includes not only visual neurons proper, but also high order,

e.g. frontal neurons that are associated with decision making, high order meta-cognition, and self-related aspects (Baars, Ramsoy & Laureys 2003).

It is important to emphasize that non-visual cortical areas as well as subcortical networks may also be involved in subjective visual experience. In fact non visual influences – such as efferent copy and attentional signals, arousal and even control of the sleep-wake cycle, are all critically important as modulatory and enabling factors – setting the excitatory level of the local neuronal assembly sufficiently high so that an ignition can occur. However, given the appropriate enabling factors – the conjecture is that the ignition-dynamics necessary for reaching a consensus state, i.e. high firing rate, sustained activity and reverberatory distribution of information can emerge completely locally in the visual assemblies proper. Unfortunately, the experimental data is still not conclusive on this point – and given its importance to the entire conjecture made here – this is where the local ignition hypothesis will stand or fall. However, although conclusive evidence is lacking – some experimental results that narrow the possibilities are already available.

The most straightforward evidence comes from the careful examination of neuropsychological cases. Thus, visual awareness appears to remain following extensive lesions of pre-frontal cortex, as well as disconnection of frontal cortex through lobotomy. In pre-frontal cortex, such lesions typically affect personality traits, decision making, self initiated acts etc. but subjects, even in cases of massive ablation or disconnection of frontal lobes appear to maintain their ability to consciously see visual images. This is in striking contrast with lesions in visual areas that lead to various and often dramatic forms of loss of visual awareness.

With regards to the more specific issue of the involvement of self-related networks in visual perception – this question has been examined in brain imaging experiments. Importantly, the results of such experiments reveal that pre-frontal "self" related cortical areas are actually *inhibited* during intense visual or auditory tasks. Activation of these areas was only apparent when subjects engaged in more reflective, introspective tasks (See Figure 4 and Goldberg, Harel & Malach 2006). These results indicate that during intense sensory experience, self-related, pre-frontal activity actually shuts off. This is probably due to the shift in the focus of attention towards assemblies related to the generation of perceptual images of the external environment, at the expense of more internally-oriented and self-related phenomenal states.

Other important structures implicated in subjective awareness are the parietal lobes. Indeed, damage to the parietal lobe in one hemisphere leads to sensory "neglect" which is undoubtedly one of the most striking cases of a selective loss of visual awareness. However, it should be noted that this dramatic effect is only apparent when the lesion is unilateral, i.e. affects one hemisphere. A bilateral lesion of parietal cortex paradoxically leads to a very different kind of deficit – the Balint
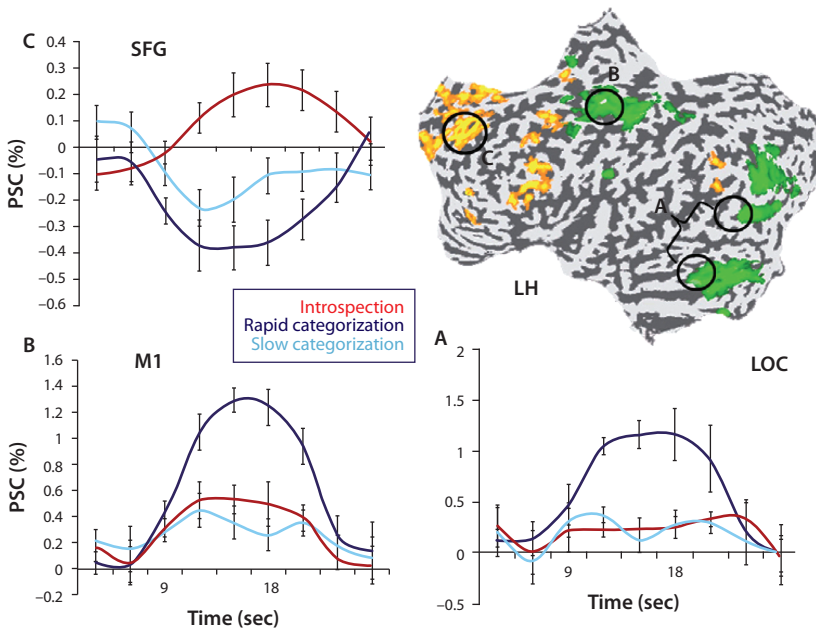
Figure 4. "Self-less" perception. Inhibition of self-related pre-frontal cortex during sensory perception. An unfolded map of the left cortical hemisphere. There is a strict segregation between self-related networks activated during introspection (orange colors) and sensory-motor networks activated during intense perceptual categorization (green). Note that activity in the superior frontal gyrus (SFG) was actually inhibited during intense perceptual categorization (SFG-inset) (Modified from Goldberg et al. 2006)

syndrome. In this syndrome, subjective awareness is largely maintained; rather, the Balint patient has difficulty in certain aspects of attentional control – such as shifting attention from one object or attending to more than one object at a time.

11. Is visual awareness dependent on higher levels of the cortical hierarchy?

When viewed in the context of the hierarchy of visual processing, it is assumed that in the visual domain this hierarchy leads to medial-temporal lobe structures such as enthorhinal cortex and hippocampus. It thus could be argued that visual awareness depends on "readout" by these yet higher order structures. Indeed, highly selective and intriguingly invariant visual responses have been recently observed in single unit recordings in such structures in human patients. Furthermore, these responses were shown to be tightly linked to subjective awareness in backward

masking paradigms (Quiroga, Mukamel, Isham, Malach & Fried 2008). However, two aspects of the responses in these regions argue against their involvement in visual awareness. First – the responses begin at long latency – about 100 msec after the ignition of neuronal responses in visual cortex – (i.e. ~300 msec after stimulus onset) which, given the rapidity of visual perception is likely too long. Second, the firing rate increases typical of such medial-temporal lobe responses are rather slow – 5–10 Hz, again arguing against a sufficiently fast information exchange that is necessitated by ignition processes (Gelbard-Sagiv, Mukamel, Harel, Malach & Fried 2008; Quiroga et al. 2008).

To conclude, the evidence with regards to the locality vs. globality of assembly ignitions is still not conclusive, although the experimental results rather constrain the manifestations of such a global spread. At this stage, this question should be viewed as a prediction of the local ignition hypothesis. Thus, if it turns out that global ignition dynamics, as conjectured above, is an essential element in any sensory percept – this would constitute a refutation of the proposed conjecture.

The proposed model is of course very qualitative and sketchy, leaving numerous open questions. Thus, it is still not clear what is the biophysical manifestation of the signals that are integrated during ignition. Are there "echoes" of the ignition that do spread out to the rest of the brain and put the ignitions in their unique context?

The ignition model emphasizes the involvement of reciprocally active neurons in the perceptual dynamics. However, it leaves open the potential role that non-active neurons may play in the generation of a visual percept. A thoughtful and quantitative model that strongly argues for a critical role for such silent neurons has been put forward by Tononi in his integrated information theory (Tononi 2005).

Another particularly thorny difficulty for the local ignition model is that it does not offer an explanation for the apparent experiential unity of subjective awareness – since, in principle at least, multiple, separate ignitions may occur simultaneously. While one could envision a global brain mechanism that prevents such multiple ignitions, the evidence for or against such control is lacking. Thus, all these questions and many more remain open. Consequently, the conjecture put forward in this chapter should be viewed more as a working framework for future research than a conclusively proven model.

References

Amir, Y., Harel, M. & Malach, R. (1993). Cortical hierarchy reflected in the organization of intrinsic connections in macaque monkey visual-cortex. *Journal of Comparative Neurology, 334*(1), 19–46.

Baars, B.J., Ramsoy, T.Z. & Laureys, S. (2003). Brain, conscious experience and the observing self. *Trends in Neurosciences, 26*(12), 671–675.

Barlow, H.B. (1972). Single units and sensation: A neuron doctrine for perceptual psychology? *Perception, 1*(4), 371–394.

Dehaene, S. & Changeux, J.P. (2011). Experimental and theoretical approaches to conscious processing. *Neuron*, *70*(2), 200–227. doi:10.1016/j.neuron.2011.03.018.

Douglas, R.J. & Martin, K.A.C. (2007). Recurrent neuronal circuits in the neocortex. *Current Biology, 17*(13), R496–R500.

Fisch, L., Privman, E., Ramot, M., Harel, M., Nir, Y., Kipervasser, S., … Malach, R. (2009). Neural "Ignition": Enhanced activation linked to perceptual awareness in human ventral stream visual cortex. *Neuron, 64*(4), 562–574. doi:10.1016/j.neuron.2009.11.001.

Gelbard-Sagiv, H., Mukamel, R., Harel, M., Malach, R. & Fried, I. (2008). Internally generated reactivation of single neurons in human hippocampus during free recall. *Science, 322*(5898), 96–101.

Gilaie-Dotan, S. & Malach, R. (2007). Sub-exemplar shape tuning in human face-related areas. *Cerebral Cortex, 17*(2), 325–338.

Goldberg, II, Harel, M. & Malach, R. (2006). When the brain loses its self: Prefrontal inactivation during sensorimotor processing. *Neuron, 50*(2), 329–339.

Grill-Spector, K., Kushnir, T., Hendler, T. & Malach, R. (2000). The dynamics of object-selective activation correlate with recognition performance in humans. *Nature Neuroscience, 3*(8), 837–843.

Grill-Spector, K. & Malach, R. (2004). The human visual cortex. *Annual Review of Neuroscience, 27*, 649–677.

Hasson, U., Hendler, T., Ben Bashat, D. & Malach, R. (2001). Vase or face? A neural correlate of shape-selective grouping processes in the human brain. *Journal of Cognitive Neuroscience, 13*(6), 744–753.

Hung, C.P., Kreiman, G., Poggio, T. & DiCarlo, J.J. (2005). Fast readout of object identity from macaque inferior temporal cortex. *Science, 310*(5749), 863–866.

Levy, I., Hasson, U. & Malach, R. (2004). One picture is worth at least a million neurons. *Current Biology, 14*(11), 996–1001.

Malach, R. (2007). The Mesurement problem in consciousness research. *Behavioral and Brain Sciences, 30*, 516–517.

Quiroga, R.Q., Mukamel, R., Isham, E.A., Malach, R. & Fried, I. (2008). Human single-neuron responses at the threshold of conscious recognition. *Proceedings Of The National Academy Of Sciences Of The United States Of America, 105*(9), 3599–3604.

Sheinberg, D.L., Leopold, D.A. & Logothetis, N.K. (1995). Controlling Perceptual Alternations During Binocular-Rivalry in Humans and Monkeys. *Investigative Ophthalmology & Visual Science, 36*(4), S668–S668.

Thorpe, S.J. (2005). Ultra-rapid scene processing – Just how fast? *Journal Of Cognitive Neuroscience*, 10–10.

Tong, F., Meng, M. & Blake, R. (2006). Neural bases of binocular rivalry. *Trends In Cognitive Sciences, 10*(11), 502–511.

Tononi, G. (2005). Consciousness information integration, and the brain. *Boundaries Of Consciousness: Neurobiology And Neuropathology, 150*, 109–126.