

a0005

Neural Basis of Perceptual Awareness

J-D Haynes, Charité – Universitätsmedizin Berlin, Berlin, Germany

© 2009 Elsevier Inc. All rights reserved.

Glossary

g0005

Binocular rivalry – When conflicting stimuli are presented to the two eyes, conscious perception can alternate spontaneously between the input to the left and the right eyes.

g0010

Choice probability – The accuracy with which an organism's perceptual choice in a decision task can be predicted from a neural signal.

g0015

Global workspace theory – A theory that postulates the neural process underlying conscious awareness is a global distribution of information throughout the brain.

g0020

Isomorphism – Literally 'identity in structure' typically refers to the notion that similarities between perceptual experiences are reflected in similarities of the underlying neural signals.

g0025

Microconsciousness – The theory that perceptual awareness depends only on suitable representations in sensory brain regions and not on additional activity in prefrontal or parietal cortex.

g0030

Perceptual threshold – When the intensity of a barely visible stimulus is gradually increased, there is no abrupt transition between 'unseen' and 'seen' but a gradual transition with an intermediate intensity range where the stimulus is sometimes seen and sometimes not.

g0035

Reversible figures – Typically this refers to visual shapes that can be seen in different ways and can give rise to different geometric or semantic interpretations.

alike. The 'qualia' or 'raw feels' of sensory experience such as the redness of red, the timbre of an instrument, or the scent of a specific flower are the most vivid aspects of consciousness. In contrast, our abstract thoughts (such as the feeling of understanding a sentence) appear to have much weaker experiential qualities.

Research on the neural correlates of perceptual consciousness has mainly focused on visual perception, which has been studied like no other field of neuroscience and especially like no other sensory modality. The visual system belongs to the best understood and most researched parts of the brain. A number of key concepts in neuroscience originated from the field of vision such as the concept of receptive field or the role of neural synchronization. The abundance of research in this field has yielded detailed mathematical models that make detailed quantitative predictions about a number of visual phenomena. Because of the rich body of research on visual perception, this article focuses primarily on the neuroscience of perceptual awareness in the visual modality. Many of these findings can be transferred to other modalities where phenomena often have direct counterparts in the visual system.

p0010

Crossing the Threshold to Awareness

A starting point for an investigation into the mechanisms of conscious visual experience could be to compare cases where stimuli are clearly visible to cases where stimuli fail to reach awareness. For example, if the intensity of a weak, invisible stimulus is gradually increased, it will at some point be strong enough to reach awareness. The intensity where the transition from 'unseen' to 'seen' occurs is called the perceptual threshold. To compare neural processing with and without awareness, one could conduct a simple experiment

s0010

p0015

s0005

Introduction

p0005

The rich qualitative properties of perception have made it a key focus of consciousness research for philosophers, neuroscientists, and psychologists

that contrasts neural responses to stimuli that are either above or below the threshold of perception. Any corresponding differences in brain activity could reflect the effect of awareness on neural processing. However, the comparison between seen and unseen stimuli would be confounded because in general the physical intensity of visible stimuli is stronger than the intensity of invisible stimuli. It would thus remain unclear whether any observed neural effects are indeed due to the difference in awareness or due to differences in the physical characteristics of the stimuli.

p0020 There are however ways to compare 'seen' and 'unseen' stimuli while at the same time avoiding stimulus confounds. When the intensity of a simple stimulus is gradually increased, there is no abrupt transition between 'unseen' and 'seen' at the perceptual threshold. Instead there is a range of intensities where the stimulus is sometimes seen and sometimes not, thus yielding a certain percentage of 'seen' and 'unseen' responses. Because there is no sharp transition from unseen to seen, the perceptual threshold is usually defined probabilistically as a specific proportion of 'seen' judgments (or alternatively a specific proportion of correct discriminations between stimulus present and stimulus absent). The gradual transition from 'unseen' to 'seen' can be used to separate physical stimulus properties from awareness in two different ways. The first approach is based on the shape of the threshold function that relates increases in physical stimulus intensity to the monotonously increasing proportion of 'seen' responses. The threshold function is s-shaped, thus visibility increases slowly for low and high physical intensities, but strongly for intermediate intensities. Because the s-shape means that visibility undergoes a nonlinear change in a range where physical stimulus properties change linearly, this allows one to separate the physics from perception by identifying brain regions that exhibit an increase in response amplitude that matches the s-shaped threshold function. Using this approach several studies have shown tight links between threshold functions for stimulus intensity and signals in early visual cortex. Threshold functions for the identification of more complex features such as objects are also closely linked to response profiles of cells in brain regions specialized in object recognition.

There is a second, even more powerful, way to study visual awareness using the gradual, probabilistic nature of perceptual thresholds. Stimuli that only reach awareness on a certain proportion of trials are very useful because they allow one to directly compare conscious and unconscious trials for the same physical stimulus parameter. For example, trial-by-trial fluctuations in perception of simple pattern stimuli are reflected by corresponding changes in activity already in primary visual cortex. Thus, already the earliest regions of the cortical visual system can closely reflect conscious visual perception of simple stimulus features. Some studies indicate that already at early stages of processing, the effects of consciousness can be stronger than the effects of physical stimulus characteristics. The differences between processing of stimuli that do or do not reach awareness is also manifest at much higher brain regions, including regions in the prefrontal cortex that are involved in top-down control of processing and in behavioral report.

A broader conceptual framework for understanding what happens in the brain when humans are viewing stimuli around the threshold to awareness is offered by perceptual decision making where subjects are required to perform simple detection and discrimination tasks. In a detection task a subject is asked to judge whether they believe to have seen a stimulus on an individual trial or not. Discrimination tasks come in two basic variants. Either the observer is shown one stimulus and has to judge which of several potential alternative stimuli it was. Or the observer is shown several stimuli and has to judge which one is which. Discrimination tasks do not probe for awareness of a stimulus but for awareness of a difference between stimuli. Performance in perceptual decision making is often accounted for by a sequence of simple information processing steps. In a first step, the presentation of a stimulus evokes a neural process encoding 'sensory evidence' about the stimulus presented. A second step consists of a 'decision variable' that is derived from the sensory evidence. The decision variable collapses all available sensory information in a way that provides for an efficient decision given the current behavioral goals. Finally, the values of the decision variable are mapped to a set of 'judgments.' For categorical

p0025

p0030

judgments such as ‘stimulus A’ versus ‘stimulus B,’ these signals reflecting the outcome of the decision are necessarily dichotomous and are typically directly related to specific motor commands that are used to indicate the judgment. Several studies have used a special type of motion stimuli to unravel the sequence of steps involved in perceptual decision making, so-called ‘random dot kinematograms.’ These are a blend of a ‘signal’ and a ‘noise’ stimulus. The signal stimulus is composed of a field of random dots coherently moving in one direction plus a noise stimulus consisting of randomly moving dots. The task of the subjects is to detect the drift direction of the coherently moving stimuli. The more the stimulus consists of signal dots and the less of noise dots the better the drift direction can be seen, thus yielding a threshold for perceptual motion detection. These stimuli allow one to calculate a so-called ‘choice probability’ that describes the accuracy with which the perceptual choice for ‘unseen’ and ‘seen’ is predicted by the activity of neurons in a specific area. A number of single-cell recordings in animals performing choice tasks with such stimuli has revealed that signals in regions of the brain specialized for motion processing partly predict the outcome of an animal’s decision – and hence presumably their perception. However, the prediction from single cells in these regions is far from perfect and choice probabilities are barely above chance. This suggests that the perceptual decision is either encoded in brain areas further downstream, or that it is encoded in pools of neurons.

p0035 Perceptual decision making models can partially explain human perceptual choices at the threshold to awareness. But additional assumptions need to be made to account for what a person consciously sees in such tasks. One question is whether the perceptual experience is more closely reflected by the sensory evidence or by the decision variable. If only a single, simple stimulus is presented, the decision variable can be equated with the sensory evidence. But tasks where multiple stimuli have to be compared require a decision variable that computes a comparison, hence a relational property. In these cases the decision variable cannot directly reflect what we see but it reflects differences between things we see. A different question is how the decision making process

relates to a person’s subjective confidence in the accuracy of their decision. Confidence in a decision is frequently taken to be a good indicator of awareness based on the notion that if we are conscious of something we know that we see it and can be confident about our judgments. But most perceptual decision making models do not treat unseen and seen conditions as qualitatively different cases where visible stimuli undergo a different processing stream than invisible stimuli. Thus, they for example do not capture the finding that visible and invisible stimuli undergo different depths of processing in the brain.

Visual Competition: Masking and Rivalry

s0015

A different way to cross the threshold between ‘unseen’ and ‘seen’ without changing the intensity of a stimulus is to render its perception difficult by introducing additional, competing stimuli. A large number of experimental approaches follow this logic. For example, in motion-induced blindness a target stimulus can pop in and out of awareness when it is presented in the vicinity of a moving set of dots. In these cases it appears as if the target fails to win the competition for awareness against the highly salient moving stimulus. A related phenomenon is flash suppression where a brief flash in the vicinity of a target (presented either to the same or the opposite eye) can strongly reduce its visibility. Flash suppression has been shown to affect processing already very early in the visual system.

p0040

Visual Masking

s0020

One of the most prominent experimental procedures for manipulating awareness is visual masking, where the visibility of a target stimulus is decreased by presenting it in close spatial and temporal proximity to a so-called ‘mask.’ For example, if a target image alone is presented for brief periods it can normally be perceived quite effortlessly. However if the brief target is immediately followed by a second image consisting of an arrangement of random lines and patterns, its visibility is strongly reduced. This phenomenon is

p0045

known as backward masking and the mask image is known as a pattern mask. Visual masking has frequently been used to study the neural correlates of consciousness. Depending on conditions, masking can lead to decreases of brain activity at various stages of the visual system. But despite being invisible, masked targets can undergo a considerable degree of processing in the visual system. Object recognition is disrupted by pattern masking only at later stages of the visual system such as the lateral occipital complex (LOC) which is specialized for recognition of objects.

p0050

There are a number of different approaches to visual masking. Masks can be effective when presented either before or after the target stimuli and can be spatially either overlapping or nonoverlapping. There are also several related phenomena, including object substitution masking and crowding. A special case is so-called metacontrast masking, where masks do not spatially overlap with the targets but share common contours with them. Metacontrast masks are most effective when there is a small time delay between the target and the mask. With very brief delays or long delays the visibility is high, thus the relationship between the target visibility and the delay follows a characteristic u-shaped function. For this reason, metacontrast masking is particularly interesting because the interaction between the target and the mask has to bridge not only space but also has to seemingly operate backwards in time. The u-shaped function can be used as a signature to identify the neural locations of masking effects in the brain. Although neural activity in V1 can be disrupted by metacontrast masking, this occurs only at later

between a conflicting stimulation of the two eyes. When conflicting stimuli are presented to the two eyes, visual perception cannot fuse them. Instead, perception alternates between seeing the stimuli presented to the left and to the right eye. There has been a long controversy regarding the exact locus of such rivalry in the brain. One dominating view is that rivalry is due to a conflict between monocular populations of cells in the early stages of the visual system. Information about the eye-of-origin is largely lost beyond primary visual cortex, and truly monocular cells can only be found in V1 and earlier stages of the visual system. In this monocular account of rivalry, the stimulus presented to one eye is suppressed from awareness because the input to this eye is attenuated at a monocular level of processing. This theory was supported by several behavioral findings. For example, the sensitivity to input from the suppressed eye is reduced, suggesting that processing of the suppressed eye is reduced. Similar support for a monocular selection was that if the input is exchanged between the dominant and suppressed eyes, perception in most cases tends to follow the eye, not the stimulus.

p0070 However, an alternative pattern-based view was subsequently suggested. This was based on findings when mixed images are presented to the two eyes that allow for perceptual grouping between the two eyes. To create such stimuli one starts with two images, say a face and a tree. The next step is to mix the two images by exchanging coherent subregions in one image by the corresponding regions in the other image. This results in two complementary images, each showing sections of a face in some regions and sections of a tree in other regions. The question is: If perception during rivalry is dominated by the input from only one eye at a time, one would expect perception to alternate between the patchy image presented to the left eye and the complementary patchy image presented to the right eye. Instead there is a tendency to see coherent percepts that combine input from both eyes into a meaningful figure. This suggests that perception during rivalry depends not only on the eye of input and thus the monocular account of rivalry cannot be the full truth. This pattern-based account of rivalry was supported by single-cell recording studies that suggested that rivalry only affects 20% of cells in regions V1,

the latest stage with substantial monocular information. Instead, the main effects of rivalry seemed to be restricted to higher levels of the visual system. In regions of the temporal lobe that are specialized in high-level object recognition, rivalry affects 90% of cells. Similarly, rivalry in humans has been shown to affect high-level regions of the visual system. For example, when rivalrous perception alternates between faces and houses, corresponding increases can be seen in functional magnetic resonance imaging (fMRI) signals measured from regions of cortex specialized for processing faces and houses. However, in the following years studies on humans have gathered evidence that rivalry might affect earlier stages of neural processing than previously believed. It has been repeatedly shown that rivalry affects fMRI signals in primary visual cortex and even has effects on subcortical processing in the lateral geniculate nucleus. The current view on the mechanisms of binocular rivalry combines the monocular and the pattern-based explanations and postulates a selection for awareness at multiple stages of the visual system. Thus, the access to awareness can be regulated at many levels of processing ranging from very early subcortical regions to prefrontal cortex.

Encoding the Contents of Consciousness

Reversible figures and binocular rivalry have long dominated experimental research on visual awareness because they allow dissociating changes in conscious perception from mere stimulation factors. When the contents of consciousness change without corresponding changes in the intensity or level of awareness of conscious perception, this might help isolate where the contents of our consciousness are stored in the brain. The rationale is that a brain region that encodes the contents of our conscious perception will need to change its activity when the contents of perception change. However, at a closer look this approach has a severe flaw. It does not allow one to disentangle brain regions that specifically encode the 'contents' of consciousness from brain regions that are 'unspecifically' involved in switching between different contents.

s0030

p0075

p0080 Besides involvement during perceptual transitions, more needs to be demonstrated to show that a brain region encodes an aspect of conscious perception. It needs to be shown that the neural responses in this brain region change in a 'content-specific' manner. The contents of consciousness can be described along a number of different dimensions. Neuropsychological data from patients with brain lesions show that different dimensions of perceptual space are encoded in a number of different brain regions. For example, there is a double dissociation between regions encoding for color and motion because awareness of both can be disrupted independently following lesions either to color-selective brain regions in the fusiform gyrus versus lesions to the motion-processing region MT. Studies on agnostic patients have revealed that the perception of complex objects can also fail independently of the perception of the simple features they are composed of. This independent drop-out of specific contents of consciousness following specific brain lesions suggests that different aspects of awareness are encoded in separable cortical regions.

p0085 Lesions can provide valuable information about brain regions encoding different aspects of awareness. But they still leave several questions open. For example, lesions to primary visual cortex cause an almost complete loss of conscious visual perception. But this does not mean that all contents of our perception are encoded in V1. Cells in primary visual cortex encode flicker at much faster rates than can be perceptually resolved. If V1 really were encoding the contents of conscious perception, then one would expect that the temporal resolution of perception would match the temporal resolution of signals in V1, which is clearly not the case. Similarly, when a stimulus is flashed into just one of the two eyes, primary visual cortex still encodes the eye-of-origin. But perceptually a subject can typically not tell which eye was stimulated, demonstrating again that V1 encodes information that does not reach awareness. Thus, these dissociations raise doubt whether signals in primary visual cortex encode contents of visual consciousness or rather serve a role of relaying information into the cortical visual system without participating in awareness directly.

p0090 Such dissociations between encoding of features in V1 and the properties of visual awareness can be

taken even a step further by examining precisely how much 'information' is encoded in a neural response and how this compares to the perceptual information available for awareness. Unperceived flicker and unconscious eye-of-origin are cases where a neural process has more information about a stimulus feature than is represented in consciousness. The complementary case would be if neural signals had less information about a stimulus than is encoded in perceptual consciousness. Take as an example color perception. For a neural population to encode color percepts, it must respond with at least one different state to each identifiable color hue. If there were fewer neural states than possible color percepts, then the neural population would not be powerful enough to encode all the possible perceptually distinguishable shades of color. This suggests a useful test: To find out whether a neural population encodes perceived color hue, one can try and decode the perceived hue from signals in that neural population. If there is indeed at least one neural signal that corresponds to each perceptual state, it should be possible to fully decode color perception from this neural signal. This argument directly addresses the content-specificity of neural correlates of consciousness. It can be used to rule out that a brain area that is involved in awareness is merely realizing unspecific enabling factors of awareness, rather than encoding specific perceptual contents.

More generally, the contents of conscious perception can be described along a large number of dimensions. For each point in the visual field, we can define its brightness, its color hue and saturation, its contrast, its speed and direction of motion, depth, and many more. Furthermore, simple features can jointly form geometric shapes and even meaningful objects. The contents of consciousness along a number of different dimensions can be jointly viewed as a complex 'perceptual space.' Using an information-theoretic framework, one can directly search for the population of neurons that encodes each specific subdimension of this perceptual space. Each subdimension of perceptual space can be viewed as a data structure that is encoded in some brain region by some parameter of brain activity.

There are several techniques available for assessing the information encoded in neural populations.

p0095

p0100

One approach is to record from 'single cells' and to relate their encoding of a stimulus with the perceptual information available to a human or animal. But this neglects the information contained in distributed networks of neurons. A more powerful approach is to take into account the full information contained in 'populations of neurons' within a brain region. For example, to understand how motion percepts are encoded in the motion-processing region MT, one has to take into account not only the activity of single neurons, but of the entire neural population. This is because even a cell that is not tuned exactly to the feature of interest can still carry information about this feature. In fact, information-theoretic analyses have shown that neurons carry only very little information about features where they show the strongest responses. This is because at the peak of the tuning curve the neural responses do not differ very much for different features. Most information is contained in the side bands of the tuning function, where the neuron changes its response rapidly with changes in physical features. Using such decoding approaches it has been shown that single neurons in regions of the medial temporal cortex that are involved in recognition and memory carry information about specific visual contents, such as thoughts about specific individuals. Despite the distributed nature of neural representation, certain cells can exhibit an incredible sparseness, meaning that each visual object is encoded by only a few cells. Such cells respond say to a picture of a specific person, but not to pictures of other people. Information-theoretic approaches to the encoding of contents of consciousness are very powerful. But, depending on the theoretical perspective, they could be considered too powerful, because they can extract information also from 'unusual' brain signals such as for example from the side bands of orientation tuning curves or from deactivations of brain activity. Some theories about awareness in contrast postulate that the contents of consciousness are encoded in the brain in the form of explicit rather than implicit representations. An explicit representation directly signals the presence of a particular feature and does not need further processing to be read out. For example, the encoding of a face in a face-selective cell constitutes an explicit representation whereas the representation of a face in the retina constitutes an implicit representation

because it requires additional processing to tell that a face is present in the spatial pattern of retinal signals.

To access the information encoded in entire neural populations, it is necessary to record from multiple cells simultaneously using so-called multielectrode grids. Such recordings can only be done in animals and occasionally in humans with implanted electrodes for diagnosis of epilepsy. A noninvasive alternative approach is offered by decoding techniques for electroencephalography (EEG), magnetoencephalography (MEG), and fMRI signals. These noninvasive techniques cannot resolve individual cells; the resolution is not even sufficient to resolve individual cortical columns (the basic units of information storage in the human brain where cells encoding similar features are clustered together across a span of approximately 0.5 mm of cortex). But these noninvasive techniques can nonetheless provide a handle on information encoded at a fine-grained scale in the visual cortex. This is possible due to small fluctuations in the topography of cortical maps that produce interference patterns that can be picked up with a standard fMRI measurement grid. Using information-theoretic decoding techniques, it is possible to access the information contained in these interference patterns, for example, about a specific type of visual feature, and to compare it to the information available to a human subject who is currently perceiving the same feature. Such decoding-based neuroimaging has been used to reveal which brain regions encode which visual features and how such information relates to the information encoded in visual awareness. This has revealed that orientation-information that does not reach awareness can nonetheless be encoded at the level of primary visual cortex, another proof of a dissociation between V1 signals and visual awareness. Because brain areas further downstream in the visual system do not encode unconscious orientation information, it is plausible that the neural correlates of the contents of visual consciousness start later in the visual system.

Interestingly, the same content-based neuroimaging techniques can be used to go even a step further in mapping conscious experience to brain processes. It is possible to investigate how relationships between elements of conscious experience are reflected in similar relationships between

p0105

p0110

brain activation patterns encoding them. Such an ‘isomorphism’ between phenomenal experience and brain activity can for example be found in the way objects are encoded in the human temporal lobe. The perceived similarity between objects is reflected in a similarity between the brain activation patterns in the object-recognition regions of the temporal lobe.

s0035 **Awareness and Specific Brain Regions**

p0115 Area V1 or ‘striate cortex’ is an excellent showcase of the different types of arguments that are put forward when discussing the potential involvement of a brain area in conscious perception. The role of primary visual cortex in visual awareness has been heavily debated. Primary visual cortex is the first cortical stage of information processing and is thus an important entry point for visual information into visual cortex. Despite the fact that other entry points exist where subcortical regions directly project to more high-level visual areas, the majority of information enters the visual system through V1.

p0120 The key role of V1 in awareness can be seen from the fact that a lesion in V1 will always cause a fully blind section or ‘scotoma’ in the visual field. The scotoma exactly matches the corresponding retinotopic location of the visual field. For example, if the upper left section of V1 is fully lesioned, this causes a blind region in the lower right quadrant of the visual field. Similar lesions in higher level visual areas typically do not cause a full dropout of visual sensitivity. They only affect the perception of specific features such as color, motion, or depth. Even lesions in V2 only affect visual acuity and contrast perception. Thus, it appears that V1 plays a special role among the visual areas in that it is necessary for visual awareness. This is supported by a number of demonstrations of close correlation between processing in primary visual cortex and visual awareness. Activity in V1 closely reflects perception of simple visual features such as contrast and brightness. Even simple forms of perceptual integration such as the perception of texture boundaries and contours can be explained from properties of V1 neurons. FMRI signals recorded in human V1 correlate with the

conscious percept during binocular rivalry. When experimental subjects train to see subtle differences between simple visual stimuli, a phenomenon called perceptual learning, the improvement in performance correlates with changes in tuning properties in primary visual cortex.

However, when arguing for a role of a specific region in visual awareness from experimental correlations caution is required. The important question needs to be addressed, whether any correlations observed indeed reflect necessary conditions for awareness or whether they are purely incidental or ‘epiphenomenal.’ The retina can also be considered a necessary condition for visual awareness because a loss of both retinae causes a complete loss in visual perception. This does not mean that the retinae are strictly necessary for awareness. Instead it means that the retina is a necessary step in the normal causal chain of events leading to conscious perception. But conscious visual perception can also be caused by bypassing the retina and directly stimulating visual cortex using implanted electrodes in patients, during surgery or by using transcranial magnetic stimulation (TMS). Also, during visual hallucinations and visual imagery, V1 is not always involved.

An alternative view would be that the disruption of visual perception following lesions to V1 simply reflects the loss of the major ‘input’ channel to the visual system. It does not directly imply that V1 is always required for every type of conscious visual experience. Furthermore, there are a number of reasons that have cast doubt on a close involvement of V1 in visual awareness. On the one hand there are theoretical reasons. For example, we can directly access, manipulate, and act upon information that is in our consciousness. To explain such access, it seems necessary to assume a direct projection from regions of the brain involved in high-level control of behavior to sensory regions encoding the contents of consciousness. However, there are no direct projections between prefrontal cortex and V1 that could support such access. A further argument can be obtained from the fact that there are striking dissociations between our conscious visual experience and encoding of information in V1. For example, we are not able to consciously tell whether a monocular stimulus is presented to the

p0125

p0130

left or right eye. But V1 encodes eye-of-origin and also many other features that fail to reach consciousness such as unconscious orientation-information and unperceived high-frequency flicker. In crowding, the visibility of a stimulus that is normally clearly visible is reduced by presenting other stimuli in the surrounding regions of the visual field. Stimuli that fail to reach awareness under crowding conditions are nonetheless processed at least up to primary visual cortex, because they leave traces of orientation-selective adaptation, suggesting that an encoding of information in V1 does not automatically lead to awareness.

p0135 Even if V1 does not encode the contents of our consciousness, several findings suggest that it might nonetheless be required for awareness. In some patients large regions of visual cortex beyond V1 can continue to respond to visual stimuli despite lesions to V1 that preclude input arriving through primary visual cortex. This means that activity in such 'extrastriate' visual regions is not sufficient in itself to produce awareness without a contribution from V1. This is further supported by the fact that although residual visual sensitivity can remain in a scotoma, the subject does not subjectively feel to be seeing anything. A subject can be above chance in guessing which stimulus was presented, but they will have the impression to be seeing nothing in that region of the visual field. Such 'blindsight' is presumably mediated by pathways into extrastriate visual cortex that bypass V1, thus lending further support for the notion that activation of extrastriate cortex can be sufficient to support visually-guided behavior but is in itself not sufficient for awareness.

p0140 The term 'extrastriate cortex' refers to a group of visually responsive brain regions beyond V1 or 'striate cortex.' An important feature of early extrastriate areas V2, V3, and V4 that they share with V1 is their retinotopic organization. This means that the topography of the visual field is largely preserved in the visual maps of V1–V4, despite undergoing a coordinate transform from Cartesian to polar coordinates. Because detailed spatial information is progressively lost at higher stages of the visual system beyond V4, this means that signals encoding the topographic spatial layout of visual perception can only be found in these early visual regions.

p0145 Regions beyond V4 exhibit an increasing specialization for different visual features. Motion is

primarily processed in a dedicated area labeled MT located in the back of the temporal lobe. The main cortical color region is located in the fusiform gyrus at the bottom of the temporal lobe, although there is some debate as to whether this is the same region as retinotopic area V4. The main cortical region for object recognition is called the LOC and is located in the lateral occipital lobe and in the posterior fusiform gyrus. Visual contents can be selectively lost from consciousness. For example, color and motion can be independently disrupted following lesions to the corresponding brain regions. After lesions to specific regions of the ventral temporal lobe, patients can lose color perception, thus perceiving the world in shades of grey while other qualities of visual perception such as motion are spared. Similarly, lesions restricted to MT cause loss of motion perception without loss in color perception. The modular encoding of different aspects of visual perception is further supported by direct cortical stimulation to patients undergoing brain surgery. Depending on which visual region is stimulated, patients report seeing complex patterns, colors, or movement.

p0150 An interesting property of cells at the higher levels of the visual system is that they respond independent of the detailed physical characteristics and context of a presented object. This phenomenon is termed 'invariance.' Say, an actress might be wearing a different dress and sporting a different hairdo, but still the cell would recognize the invariant person encoded in the image. But this invariance also comes at a price. If a cell responds to a complex visual object independent of the constituent features, it means the cell has lost all information about the fine-grained features such as brightnesses and colors that the object is made up of. This means that the encoding of contents of consciousness necessarily occurs at multiple levels. Brain regions coding complex, invariant aspects of our visual experiences can in principle not be coding the simple features. The fact that the processes underlying consciousness are fragmented, modular and multilevel is well documented, but this stands in direct contradiction to our impression that visual experience is unified rather than divided into a number of different features. The distributed encoding of the contents of consciousness in extrastriate visual areas requires additional

assumptions in order to explain the unity of consciousness. This problem is known as the ‘binding’ problem. The most prominent but also controversial explanation of binding assumes that features are bound by synchronization of cells in different brain regions.

p0155 An important dissociation in the visual system is that between representations for action and representations for awareness. There are two major visual pathways in the extrastriate visual system. The dorsal pathway is involved in spatial transformations and actions, whereas the ventral pathway is involved in object recognition. Patients with lesions of the ventral pathway cannot describe a visual stimulus any more, but can still perform visually-guided actions to the input. Patients with lesions of the dorsal pathway show the opposite pattern of disorders. They cannot perform visually-guided behavior but can consciously describe the stimulus. This has been interpreted to mean that only the ventral visual pathway supports conscious perception, whereas the dorsal pathway supports unconscious visual guidance of behavior, similar to the blindsight that follows lesions to primary visual cortex.

p0160 Supramodal regions beyond the visual system also play an important role in perceptual awareness. When a stimulus crosses the threshold to awareness, this also leads to changes in activation of regions of prefrontal cortex. For example, during studies where subjects are required to recognize masked versus unmasked words, conscious perception is correlated with increased activation in several prefrontal brain regions. Similarly, during perception in binocular rivalry, a frontoparietal network is involved in perceptual transitions, suggesting that these regions might be involved in awareness. One possibility is that these prefrontal brain signals reflect the global distribution of information as postulated by the so-called global workspace theory of consciousness. According to this theory, neural representations reach awareness when they are distributed to other brain regions. An alternative view is that content-specific brain activity within individual sensory brain regions is sufficient for awareness of these contents to occur, a theory referred to as microconsciousness. To date it is unclear whether the frontal and parietal processes involved in awareness are indeed content-specific and thus can be claimed to reflect

a distribution of information. Instead they could reflect unspecific processes, for example, they might be involved in causing awareness of a stimulus without themselves encoding the contents of consciousness. Alternatively they could be involved in the subject consciously noticing and reporting a change in the contents of their consciousness.

Another important question for the neural mechanisms of awareness is whether they involve large-scale integrative processes that jointly involve a number of brain regions. One line of thought posits that awareness of representations is related to specific dynamic interactions between populations of neurons. For example, one controversial theory suggests that awareness and binding are both closely related to synchronization between the distributed population of cells encoding the different features of an object. A similar hypothesis states that awareness of representations is related to the stage of recurrent cortical processing where feedforward and feedback neural processes overlap in early sensory regions. This is supported by several invasive studies in monkey visual cortex, as well as by studies using TMS to knock out later stages of processing in V1. When MT, the main cortical motion-processing region, is stimulated using TMS, this can create the illusory perception of movement. When a second TMS pulse is administered to V1 just following the first pulse to MT, the perception of movement is abolished. This could mean that awareness of motion depends on intact feedback projections into V1. In contradiction to this finding, patients with the so-called Riddoch syndrome who have lesions to V1 can have selectively spared motion perception despite global blindness. Thus, recurrent processing involving V1 activity cannot be a strictly necessary condition for visual awareness of motion.

Attention and Awareness

Awareness and attention are so closely related that many researchers consider them to reflect the same process. A key feature of both attention and awareness is their selectivity. Only a small amount of sensory information that is available at a given time reaches awareness. This limitation is most clearly apparent when there is a lot of competition for access to our consciousness. For example in a

p0165

s0040

p0170

crowded room we might fail to immediately notice a friend who is present although he is clearly visible. In such situations we have to consciously scan our environment by focusing our attention on one person at a time, and only when attention is focused on a known person we will be able to recognize them. Such cluttered scenes help understand the neural mechanisms that underlie the selectivity of visual awareness. According to the biased competition model of visual attention, stimuli are in competition for access to processing resources. When several stimuli are present in the visual field (or in the receptive field of cell), attention is required to bias or boost the processing of a selected visual feature, and the unselected stimuli are suppressed from further processing. Such effects have been demonstrated for single cells as well as for population responses measured with fMRI. The general finding is that competition increases at advanced stages of visual processing where receptive fields increase in size and competition occurs between an increasing number of elements in the visual field. This model of competitive interactions between processing units provides a model for understanding the selectivity of conscious visual experience.

p0175

The selectivity of attention and awareness seems to imply that the two might be two aspects of the same process. In this view attention functions as a gatekeeper to consciousness. There are many demonstrations of the tight relationship between attention and awareness. Lesions to attentional control regions can cause a disorder of awareness known as visual hemineglect, where stimuli can fail to be noticed when presented in the contralesional visual field. Another striking example is a phenomenon called inattentional blindness, where subjects fail to notice highly salient events in their visual field when their attention is directed elsewhere, thus highlighting the importance of attention for visual awareness. A related phenomenon is change-blindness, where subjects fail to notice marked changes occurring in their visual environment. In the typical experiment, a person is monitoring a cluttered visual image that is repeatedly flashing on and off for any changes made between two successive presentations. Most people have the intuitive assumption that they perceive the entire visual scene surrounding them and thus would definitely notice a salient change. In contrast to this intuition,

most people fail to perceive quite marked changes unless explicitly paying attention to the particular region of the visual scene where the change occurs. These experiments can be useful for tracking down the neural correlates of consciousness. Major changes to visual scenes that go unnoticed can still continue to be registered and analyzed up to higher stages of the visual system. fMRI signals from high-level object-selective regions signal changes in a visual scene that are not registered. Similar findings have been obtained in single-cell recordings in human medial temporal areas related to object recognition and memory.

Despite the close link between attention and awareness, there are also several findings that demonstrate that the two cannot be fully equated, although they are clearly closely related. One important question is whether stimuli outside the focus of our attention can really not reach our conscious awareness. There are several experiments that suggest that this might be the case. In one line of experiments, the attentional resources of subjects are engaged maximally at one region of the visual field, typically the center of gaze, by requiring them to perform a very difficult visual discrimination task. Despite the full engagement of attention, they can still continue to perceive certain stimuli presented at a different region of the visual field. Simple, salient stimuli such as colors or shapes can be identified without attention. Even more complex aspects of a visual scene can be perceived in parallel without requiring attention. For example, observers can readily and rapidly perceive whether an image contains an animal or not, even when the animal is hidden in a cluttered natural landscape. EEG signals recorded from humans during such tasks show signatures of rapid recognition as early as 120 ms after the onset of a stimulus. In contrast, complex combinations of geometric features cannot be perceived rapidly while attention is directed elsewhere. This suggests that the ambient 'fringe' surrounding the focus of our attention can still enjoy a fair degree of complex, high-level cortical processing and lead to awareness.

p0180

Summary

A large number of experimental studies have contributed strongly to our understanding of the

s0045

p0185

neural mechanisms of visual perceptual experience. The contents of perceptual experiences are encoded in distributed subregions of modality-specific cortex. The access to consciousness appears to be regulated at multiple stages of processing, reaching from early subcortical processing to high-level regions involving prefrontal and parietal cortex. However, several theoretical debates regarding the specific processes involved still await further clarification. It is currently unclear whether representations that reach awareness are globally made available for further processing across the brain. And it is also unclear how the many subdimensions of perceptual space that are distributed across many brain regions give rise to a unified perceptual experience.

ELSEVIER SECOND PROOF