

**Why cognitive penetration of our perceptual experience
is still the most plausible account**

Albert Newen^a, Petra Vetter^b

^a Institute for Philosophy II, Ruhr-Universität Bochum, Universitätsstraße 150, 44780
Bochum, Germany – E-Mail: albert.newen@rub.de

^b Department of Psychology & Center for Neural Science, New York University, 6
Washington Place, New York, NY 10003, USA – E-Mail: petra.vetter@nyu.edu

Corresponding author: Albert Newen – albert.newen@rub.de

Why cognitive penetration of our perceptual experience is still the most plausible account

Albert Newen, Petra Vetter

Abstract: To what extent is our perceptual experience influenced by higher cognitive phenomena like beliefs, desires, concepts, templates? Given recent arguments against the possibility of cognitive penetration, we present striking evidence against the impenetrability claims. The weak impenetrability claim cannot account for (1) extensive structural feedback organization of the brain, (2) temporally very early feedback loops and (3) functional top-down processes modulating early visual processes by category-specific information. The strong impenetrability claim could incorporate these data by widening the “perceptual module” such that it includes rich but still internal processing in a very large perceptual module. We argue that this latter view leads to an implausible version of a module. Therefore, we have to accept cognitive penetration of our perceptual experience as the best theoretical account so far given the available empirical evidence. We outline that this does not have any problematic consequences for the relation between perception and cognition.

Keywords: (max. 10)

Cognitive penetration

Impenetrability

Visual perception

Early visual processes

V1

Top-down modulation

Predictive coding

Perception–cognition boundary

Philosophy of mind

Perceptual learning

1 Clarifying the claims and shifting the burden of proof

What is cognitive penetration? Cognitive penetrability is a phenomenon that occurs if higher-level cognitive phenomena (states or events or processes), such as beliefs, desires, intentions, concepts etc., directly influence our perceptual experience. In other words, if cognitive penetration takes place, what one believes, desires, intends, etc., may alter what one sees, hears, etc. To get the debate adequately off the ground, we need to be more precise. As a first step, we can rely on the definition of the necessary conditions for cognitive penetration (being altogether sufficient) as offered by MacPherson (2012). What is held constant is: 1. the object or scenario causing the visual input, 2. the perceptual conditions, 3. normally functioning sensory organs and 4. the absence of spatial attentional shifts. If a higher-level cognitive process (we will speak of a “process” as the paradigmatic phenomenon but also allow for states or events) can nevertheless change the perceptual experience, this is a case of cognitive penetration in a narrow sense. What penetrates and what gets penetrated? Just to repeat: the penetration comes from an activation of higher cognitive processes like beliefs, desires or concepts, more precisely, it comes from the activation of the content of these higher cognitive processes. The object of penetration is my perceptual experience. What is the nature of the penetrating relation? Regarding the question on whether cognitive penetration is a causal (Siegel, 2005) or a rational (Pylyshyn, 1999) relation, we take the side of Siegel and presuppose that it is sufficient if there is a *relevant* causal relation¹ between the content of a higher cognitive process and the percept that is influenced by it (given that the content of the higher-cognitive process is part of the relevant causal factors). On the contrary, the claim of cognitive impenetrability is the principle claim that cognitive penetration of our perceptual experience never happens. So far, the community is still concerned with the principle debate of whether perceptual experience is penetrable or impenetrable.

Before entering the discussion, it helps to clarify what exactly the claims of the two opposing camps in the discussion are. In the classic paper of Pylyshyn (1999), he defends a weak version of cognitive impenetrability since he limits his claim such that cognitive impenetrability is only valid for early visual processes. Thus, he allows for two possible interventions of higher cognitive processes onto perceptual processes that do not conflict with an impenetrability thesis: 1. the modification of sensory input by change of attention prior to early visual processes and

¹ It is common ground in the literature that we need to exclude some causal chains as irrelevant, e.g. if my thinking about an upcoming exam causes migraine and this induces light flashes as part of my visual experience. This would not count as a case of cognitive penetration as the light flashes are not systematically caused by my thinking about the exam, but are only accidentally related. For most persons, thinking about an exam does neither cause migraine nor does it lead to light flashes. To exclude these cases, some philosophers argue that we need to presuppose a rational (or semantic) relation in addition to a causal chain. But the latter is too demanding (without further argument) since it is more straight forward to demand a *relevant* causal chain. Furthermore, let us presuppose – for the sake of argument – that the following would be a paradigmatic case of cognitive penetration, e.g. we activate different color concepts and thereby have a different color experience (while all the other perceptual conditions as defined above remain fixed). It seems to be going much too far to claim that we need a rational or semantic relation concerning this effect: Cognitive penetration just claims that there is a modified perceptual experience caused in a *relevant way*. Why should we presuppose that any step in the complex causal route resulting in a modified perceptual experience must not only be a causal but also a rational or semantic relation (as argued by MacPherson, 2016)? Why should we exclude such an ideal case if there is a causal chain involved which we do not understand any more than just by the fact that it is a non-accidental and systematic way the brain processes such situations?

2. the “interpretation” which takes place after early visual processes but – as we read him – prior to the completion of our perceptual experience.

„Our hypothesis is that cognition intervenes in determining the nature of perception at only two loci. (...) These two loci are: (a) in the allocation of attention to certain locations or certain properties prior to the operation of early vision (...) (b) in the decisions involved in recognizing and identifying patterns after the operation of early vision. Such a stage may (or in some cases must) access background knowledge as it pertains to the interpretation of a particular stimulus.” (Pylyshyn, 1999)

Thus, Pylyshyn defends cognitive impenetrability only for a module of early visual processes (leaving those underdetermined). In an updated view, Raftopolous (2014) even constrained an early visual processing module to those areas which are involved in the first 100 msec after visual stimulation.

A much more radical claim is taken in the recent paper by Firestone and Scholl (2015). After having discussed important methodological pitfalls and thereby setting a new bar for proofs of cognitive penetration, they express their view very clearly: “Until this high bar is met, it will remain eminently plausible that there are no top-down effects of cognition on perception.”

Thus, they claim that influences from higher cognitive processes only occur prior to any kind of visual processes and after the completion of a visual percept: in between, visual perception is created in a large visual processing module that is cognitively impenetrable.

What is our strategy of argumentation? To oppose the claim of cognitive impenetrability, we should consider at least the two versions of the claim: the *weak impenetrability claim* that only some early visual processes form a cognitively impenetrable module (Pylyshyn, 1999) or the *strong impenetrability claim* that all processes forming our visual experience are cognitively impenetrable (Firestone & Scholl, 2015); in both cases the perceptual conditions mentioned in the definition above should remain constant, i.e. the visual input, the focussed attention, normally functioning sensory organs as well as the external perceptual conditions. In our defence of cognitive penetrability we will present the most striking evidence and aim for a scientific evaluation of the most plausible theory accounting for it. We will conclude that despite lacking an *experimentum cruxis* for the time being, there is a large amount of evidence supporting the thesis of cognitive penetration according to the principle of best explanation. This line of approach also requires highlighting the most important evidence in favour of the opposing impenetrability thesis. However, our argumentation will shift the burden of proof towards the defenders of cognitive impenetrability. Furthermore, we will outline the most plausible view of the interface between perception and cognition that is compatible with our main defence of cognitive penetrability of perceptual experience.

Why should we accept the radical claim of cognitive impenetrability in the first place? What is the most striking evidence in favour of cognitive impenetrability? First of all, we have to account for the phenomenon of visual illusions that remain stable even if we are fully informed about the illusionary status of our experience. Let us stick with the simple and well-known example of the Müller-Lyer illusion: two arrows (with different orientation of the arrowheads) are perceived as having different length even after measuring the length of the main lines and assuring myself that they are indeed of the same length. My knowledge does not influence my perceptual experience in such a scenario. However, it does not follow from this observation that

all perceptual experiences in all perceptual scenarios can never be penetrated by contents of higher cognitive processes. The reply of the opponent could be that visual illusions allow us to discover construction principles of perception which remain active in everyday perception as well. Again, we agree that these construction principles exist and that they remain active but we disagree that that these construction principles can never be influenced by higher cognitive processes just because they remain dominant or compete with each other in the special perceptual circumstances of visual illusions. Generalising from the special case of visual illusion to everyday visual experience is a radical overgeneralization which lacks supporting evidence. We can grant that there are basic principles of constructing a perceptual experience of line length (or orientation etc.) which cannot be influenced by our beliefs in many situations. But why should we accept that they can in principle *never* be influenced by *any* higher cognitive content? This must at least presuppose that visual illusions are the standard cases of perception but in fact the contrary seems to be the case: if visual illusions are the rare exceptions of our everyday visual experiences, we must be careful to generalize from the exceptional cases to the standard cases without any further evidence. The second main argument in favour of cognitive impenetrability is the functional specialization of visual brain areas: if there is a lesion in color processing area V4, we perceive the world in shades of black and white lacking any colors. If someone suffers from a lesion in motion area V5, she is unable to perceive the motion flow of filling a glass with water. Functional specialization of visual brain areas is a major discovery. But what does it entail? At most we can argue that some type of information (color information) can only be adequately processed to become part of our perceptual experience if it is processed by this module. Maybe V4 is domain-specific in the sense that it mainly processes color information but it does not follow at all that V4 cannot be influenced by higher cognitive contents. Functional specialization of brain areas by itself does not imply cognitive impenetrability. We suspect that this unjustified implication is based on the fact that Fodor (1983) defined modules with several main criteria combining domain-specificity, impenetrability and being innate. But this definition of combined module criteria should not mislead us. Without further evidence the definition is just not well chosen: domain-specificity and impenetrability need not go together. To come to an intermediate conclusion: We are lacking convincing evidence for the thesis of cognitive impenetrability of our perceptual experience in the first place and thus the burden of proof lies on the proponents of this claim.

We will now describe positive evidence supporting the thesis that cognitive penetration of our perceptual experience can in fact take place. The positive evidence is sometimes also used to make an argument against one or both versions of the impenetrability thesis. Since Pylyshyn constrains his claim to early visual processes together with the idea that these early visual processes are implemented in an impenetrable perceptual module, this claim can be discussed in light of recent empirical findings concerning perceptual processes and brain modularity. In the case of the strong impenetrability thesis (e.g. Firestone and Scholl, 2015), it is more difficult to reject (or defend) an impenetrable perceptual module: although a vision module is usually defined, it is not constrained with regard to what it should contain in terms of functional brain areas. The observation of stable visual illusions is often taken as a main argument for impenetrability. However, visual illusions only indicate that in some exceptional perceptual scenarios there is a difference between some low-level processes and our knowledge that is unable to penetrate the illusion. As a consequence, a very large unspecified black box is

described containing all processes after the sensory stimulation has happened and before the creation of the percept is completed. This view entails several central problems. To repeat one worry: Why should we accept the generalization of exceptional cases of visual illusions to any case of everyday perceptual experience? Further conceptual problems are: If the black box contains such a large amount of processes, why should we believe that this black box does not contain types of processes which, under non-illusory circumstances, clearly count as higher cognitive processes? Maybe this large black box could involve typical cognitive processes right from the start? It may be possible that in some perceptual scenarios, low-level processes dominate higher cognitive processes in the creation of the visual percept. Finally, such a big black-box-module is not functionally specialized and contains a great variety of processes. This makes it more difficult to prove that it nevertheless should remain impenetrable. Thus, the larger the supposed vision module, the higher the burden of proof for the impenetrability claim to be plausible.

A similar argument can be made for the weak impenetrability thesis. If one supposes some basic and very early visual processes as being impenetrable, then this asks for a module that becomes smaller and earlier every time research methodologies achieve a higher resolution. The principle problem lies in the postulation of a clear-cut and impenetrable module *per se*. As we argued above, while some brain areas are certainly functionally specialised to some predominant functions, this does not mean they have clear-cut functional boundaries that cannot be penetrated by contents from other brain areas. Postulating an impenetrable module, whether implausibly large or very small, will always entail a boundary that is not tenable. We will provide further arguments against such a module in the following section.

2. Arguments against an encapsulated and impenetrable perceptual module

2.1. Structural arguments

Both from a functional and neuroanatomical perspective, there is no evidence for an encapsulated and impenetrable visual module. Of course, functional specialisation in the brain exists, for example, V4 predominantly processes color, and V5 predominantly processes motion, but this does not exclude the possibility that these areas process, to a weaker extent, also other types of information. For example, in the case of the fusiform face area (FFA), it was initially argued that this area only processes visual information from faces (Kanwisher et al., 1999; Yovel & Kanwisher, 2006) until, later on, the FFA was shown to selectively respond also to cars, birds and other object categories (Gauthier et al., 2000; Cukur et al., 2013). This does not deny that FFA predominantly processes faces, the novel evidence just denies that it uniquely processes faces. Likewise for many other specialised brain areas, it is well possible that the future will reveal additional functions that we cannot even imagine yet. In the case of V1, despite it is being one of the most extensively studied area in the brain, models can so far explain only up to 40% of V1's processing variance during natural vision (David & Gallant, 2005; Carandini et al., 2005). That is, there is a lot of remaining unexplained processing function in V1 that may surprise us in the future. Thus, assigning a sole function to a specialised brain area is a risky endeavour as we simply do not know enough about even a single brain area. It is therefore unjustified to speak of a brain area as restricted to *one* functional role. Thus, against

the traditional view of Pylyshyn (1999) and Raftopolos (2014), we simply lack evidence for a perceptual module in the strict sense of a functionally isolated module. And even if we grant the existence of a somewhat functionally specialized module, this still does not decide on its being cognitively impenetrable or not. Let us now have a brief look on the question of impenetrability from the structural perspective.

Also from a neuroanatomical perspective it is unjustified to speak of impenetrable modules in the brain. Recent evidence shows that cortical brain areas are much more heavily interconnected than previously thought, namely to 66%. That is, each brain area is connected to 66% of the rest of the brain (Markov et al., 2013). While the majority of these connections are short-distance connections between neighbouring areas, the hierarchical structure of the brain implies a cascade of interconnected brain areas across processing levels. As Gilbert and Li (2013) put it:

“The source of top-down influences can be widespread, either by direct connections from different cortical areas or by a cascade of inputs originating from many more areas. In effect, a large part of the cerebral cortex can exert influences over individual neurons within a particular area, with multiple descending inputs interacting with intrinsic cortical connections. As such, each neuron is a microcosm of the brain as a whole, with synapses carrying information originating from far flung brain regions.“

Both Markov et al., (2013) and Gilbert and Li (2013) base their claims on a vast amount of empirical evidence (most of it being unknown at the time Pylyshyn wrote his 1999 article, and entirely ignored by Firestone and Scholl, 2015).

That is, given that each brain area is highly interconnected to a large amount of other brain areas, it is implausible to suggest that any brain area is strictly specialised and impenetrable by other brain areas including higher cognitive processes. Thus, from a structural point of view, penetration of a visual module can easily be realized by higher cognitive areas via a top-down cascade of intermediate brain areas. All these connections in the brain have evolved throughout human and mammalian evolution, and are pruned during development, they must serve a function and are not just an epiphenomenon. Therefore, acknowledging the anatomical fact of a highly interconnected brain makes an impenetrable brain module implausible.

Particularly for the case of vision, recent years have seen a large amount of empirical evidence for feedback connections to visual cortex and their role in exerting top-down influences. Many other authors have reviewed this evidence in detail (e.g. Gilbert & Li, 2013; Petro, Vizioli & Muckli, 2014). While the exact role of feedback connections is under-researched, e.g. with respect to the nature of the information they exactly convey and how exactly they are involved in visual processing, their structural and functional existence is widely accepted. Again, from this structural and evolutionary point of view, it is implausible to assume that all this large amount of feedback connections to visual cortex is a mere epiphenomenon without functional significance in brain processing and creating perceptual experience. In light of all this widely accepted evidence, the claim of a functionally encapsulated and impenetrable perceptual module is simply not plausible, neither for the weak nor for the strong impenetrability claim. Thus, concerning the structural organization of the brain the burden of proof lies on the proponents of an impenetrable and encapsulated perceptual module to show where and how it could be implemented in the brain.

As we will also see in the following section, the bidirectional structure of feedforward and feedback connections mediates constant and recurrent feedforward and feedback processing that is necessary for conscious perceptual experience (e.g. Lamme & Roelfsema, 2000). This recurrent feedforward and feedback processing seems to be particularly important for naturalistic complex visual stimuli that require segmenting objects in the foreground from a cluttered background akin to our natural visual circumstances (e.g. Hupe et al., 1998; Scholte et al., 2008). Many visual phenomena are studied with simple and impoverished visual stimuli that subchallenge the visual system. Also some visual illusions that have been used to support the impenetrability claim are rather impoverished stimuli (e.g. the Mueller-Lyer illusion). It is of course easy to propose an impenetrable perceptual module with a single feed-forward sweep for very simple and impoverished visual stimuli because these stimuli do not tax the visual system very much. However, this model is not functional in naturalistic vision – for visual perception in a complex and cluttered environment as we are confronted with in everyday life, recurrent processing between many visual brain areas is necessary and no meaningful perceptual content could be derived from a single impenetrable module.

2.2. Arguments from temporal processing

Sometimes an early vision module is proposed on the basis of some basic visual processing supposedly happening very fast and bottom-up, too fast to be affected by top-down processing. However, the original notion of “early” and “late” processes has been challenged by evidence showing that even supposedly “late”, higher level processes occur much earlier than previously thought. Time-resolving electrophysiological evidence showed that visual cortex is activated within 50 ms and pre-frontal areas within 80 ms after visual stimulus onset. This leaves plenty of time for iterative top-down processing between “cognitive”, e.g. frontal and parietal, areas and sensory, e.g. occipital, areas, within the first 100 – 200 ms after visual stimulation (Foxy & Simpson, 2002). Thus, complex high level and reiterative processing can happen very fast and can influence visual processing very early on (Michel, Seek & Murray, 2004; Plomp et al., 2015).

Within the visual system, there is considerable evidence for fast top-down processing already within the first 50 ms after stimulus onset, certainly between motion area V5 and primary visual cortex V1 during motion perception (e.g. Pascual-Leone & Walsh, 2001; Silvanto et al., 2005; Vetter et al., 2015), sometimes even shown to occur as early as 10 ms (Hupé et al., 2001). In turn, the frontal eye fields (FEF), a higher-level area in frontal cortex involved in motor planning of eye movements, exerts its influence to V5 within 30 ms (Silvanto, Lavie & Walsh, 2006). Therefore, a feedback loop from a frontal region to an early occipital region can take as little as 80 ms or less. Importantly, this feedback happens in a task-specific manner, telling us something about the information conveyed in this feedback: when the task requires face recognition, FEF signals are sent to face-sensitive regions and when the task requires motion discrimination, FEF signals are sent to motion area V5, both within a time frame of 20-40 ms after FEF activity (Morishima et al., 2007). Therefore, top-down signals from a frontal region to earlier visual regions do not occur in a general and unspecific manner. Instead, they carry task-specific, and thus higher-level cognitive information, and are transmitted very quickly.

These fast top-down signals from FEF have actual perceptual consequences: For example, during visual conjunction search (searching for an item defined by the conjunction of two features, e.g. a red horizontal line in a display of green and red horizontal and vertical lines), FEF exerts its perceptual effect as fast as 40-80 ms after visual stimulus onset (O’Shea et al., 2006). Also object shape recognition involves recurrent “round-trip” feedforward and feedback processing within 60 ms (Drewes et al., 2016). Now, whether object shape recognition or conjunction search is classified as a cognitive or purely perceptual process is a distinction that is key to the debate on cognitive penetrability but hardly addressed. We will come back to this question later. However, in many discussions on cognitive penetration, object shape recognition has been considered to be a cognitive process.

What follows from this (non-exhaustively reviewed) literature is that top-down influences from higher level areas to “early” visual cortex can occur very early, and much earlier than previously thought. Thus, it is unjustified to conclude from an early timing of a visual process for it to be unaffected by top-down influences or feedback. In fact, it is more likely that both feedforward and feedback processes happen in parallel rather than in a strictly serial manner. Hupe et al. (2001) say: “[...] Cortico-cortical connections in the visual cortex must be conceptualized as a network of interacting areas responding with near-simultaneity, rather than as a pipeline-type architecture.” All this empirical evidence especially speaks against the weak impenetrability claim in a version where early visual processes are defined by temporal criteria as e.g. proposed by Raftopoulos (2011, 18): “Early vision includes a feed forward sweep (FFS) in which signals are transmitted bottom-up. In visual areas (from LGN to IT) FFS lasts for about 100 ms.” Again, the burden of proof concerning temporal processing is shifted towards the proponents of a weak impenetrability claim: they must demonstrate that there still remains a dimension of visual processing that is never influenced by any top-down signals originating in cognitive processing areas. This line of argument runs into the trap of postulating an implausibly small impenetrable “perceptual” module as already mentioned at the end of section 1. Even if the existence of a very small module that is extremely constrained in spatial and temporal processing could be demonstrated, this would no longer be a plausible candidate for a *perceptual* module. Whatever the function of such a small module, it would no longer be sufficient to constitute the core part of perception (given the evidence we have already presented).

2.3. Functional evidence for categorical top-down influences to early vision: Category-specific information can be transferred to early visual cortex

Let us pick one example from our own work demonstrating how feedback from higher brain areas, carrying categorical and semantic information, affects early visual cortex, thus providing evidence against an impenetrable early vision module. Vetter et al. (2014) presented human participants with natural sounds, e.g. sounds from naturalistic environments such as people talking at a party, traffic noise from a street scene and bird singing from a forest scene. Participants listened to these sounds while lying in an MRI scanner and being blindfolded, that is, there was no visual information entering the brain. Thus, while visual cortex was never actually stimulated with feed-forward retinal input, Vetter et al. (2014) showed that neural activity patterns in early visual cortex were distinct depending on the semantic content of the

sound. That is, different natural sounds elicited distinct neural activity patterns in early visual cortex in the absence of visual stimulation. Given that there was no feedforward visual stimulation, the content-specific sound information must have been transferred to early visual cortex by feedback from other parts of the brain.

If one wanted to interpret these results in agreement with an impenetrable early visual module, one would have to make several assumptions. For example, one could argue that auditory perception of natural sounds does not require much cognitive processing and can be achieved within an early perceptual module. Then the auditory information is passed directly to vision within the same perceptual module, ending up as distinct neural activity patterns in early visual cortex. This interpretation would be in line with a possibly impenetrable perceptual module. However, this interpretation is implausible for many reasons. First, it assumes that early audition and early vision happen on the same early perceptual analysis level and thus in the same perceptual module. However, this argumentation entails that audition and vision are part of the same perceptual module violating an important criteria of a module, namely its functional specialisation. Sensory specificity of brain areas has always served as one of the main arguments for functional specialisation and modularity: the fact that visual cortex is specialised for visual input and auditory cortex for auditory input. Merging both vision and audition in one perceptual module violates one of the most basic criteria for a module. Second, this interpretation assumes that complex natural sound analysis and recognition can be achieved within the perceptual module. However, complex sound analysis requires auditory association areas extending beyond primary auditory cortex (e.g. Schirmer et al., 2012, for review). Thus, this model is only plausible if one assumes a relatively large perceptual module encompassing complex semantic sound analysis. Assuming such a big module runs again into the same problem as mentioned above: to what extent can a perceptual module be called purely perceptual and not cognitive, if it encompasses complex semantic analysis? We will come back to this point again later on. Third, such a perceptual module must be realised in the brain such that audition and vision can communicate directly with each other without involving higher level, extra-modular brain areas. Anatomically speaking, there are direct anatomical connections between early auditory and early visual cortex (Eckert et al., 2008; Beer et al., 2011), but they are very sparse in number. If one wanted to assume that these connections are the only way auditory information is communicated to early visual cortex, then sound information should be distinguishable only in early auditory and only in early visual cortex. A whole-brain analysis performed in Vetter et al. (2014) demonstrated that this is not the case: sound content could also be distinguished in a large part of auditory association cortex and in multisensory brain areas such as posterior superior temporal sulcus and precuneus. If one wanted to assume an intra-perceptual interpretation, the perceptual module would have to be very large and comprise all these higher-level auditory, visual and multi-sensory areas.

There are further arguments from Vetter et al. (2014) that speak against an impenetrable module explanation. One of the additional experiments showed that sound content can be distinguished in early visual cortex even when there was neither auditory nor visual stimulation, but when participants merely imagined the sounds. That is, even auditory imagination (induced by the participants' will upon a word cue), including auditory memory recollection, induces content-specific activity patterns in early visual cortex. Furthermore, it was shown that the information

transmitted to early visual cortex was shared when sounds were actually heard or when they were merely imagined. Thus, also the semantic content of mental imagery is transferred to early visual cortex. Again, to explain these results within an impenetrable module, one would have to assume a very large module that encompasses mental imagery and its semantic information content without involving any cognitive processing. Such an assumption makes the notion of a perception-cognition boundary very imprecise and arbitrary.

Furthermore, there is another crucial result of the Vetter et al. (2014) study: the information from sounds that ends up in early visual cortex is not only content-specific, but also category-specific. That is, when participants are presented with sounds from different scenes that belong to the same semantic categories (e.g. traffic noise and the sound of a starting airplane, belonging to the category of inanimate sounds), it is the information shared among sound exemplars within this semantic category that is transferred to early visual cortex, not the information specific to the features of a particular sound exemplar. That means the information that ends up in early visual cortex is categorical and high-level, and thus very likely of semantic nature. This also entails that the decoded top-down information does not merely stem from a reactivated mental image being similar to, or a weaker version of, a representation induced by retinal feedforward stimulation (Pearson et al., 2015). Such a weaker reactivation of a visual representation could in principle still be part of a big perceptual module, but the fact that abstract and categorical information is represented in early visual cortex speaks for a higher-level origin and excludes a low-level feature driven mental imagery influence. Explaining these effects as all occurring within a single impenetrable perceptual module would imply a very large perceptual module processing auditory and visual information, semantic analysis and categorisation and mental imagery. Such a module is conceptually implausible as it would bundle a large amount of specialised brain regions and could hardly be defined as purely perceptual anymore.

The alternative and more plausible explanation is that semantic sound content is extracted in auditory association cortex and transmitted through a cascade of feedback connections via multi-sensory (and likely imagery-related) brain regions, such as the posterior superior temporal sulcus and the precuneus, all the way down to early visual cortex. This explanation does not involve giving up functional specialisation, quite to the contrary: brain areas maintain their functional specialisation (e.g. for auditory or visual specificity), but are permeable to input from other brain regions, both higher up or lower down in the hierarchy, and sending recurrent signals back and forth. Like this, no large module has to be defined or an arbitrary distinction between perceptual and cognitive brain areas be drawn. The idea of functionally specialised but penetrable brain areas being in constant recurrent communication with other brain areas is in line with a large amount of functional and anatomical evidence as mentioned above.

It should be noted that the study by Vetter et al. demonstrates semantic top-down influences to early visual cortex, but does not directly show influence on visual perception *per se*. On the behavioural level, cross-modal priming studies demonstrated that semantically congruent sounds improve sensitivity of picture identification (Chen & Spence, 2011a &b) and facilitate conscious access to a matching image (Chen, Yeh & Spence, 2011). On the neural level, however, there is so far little evidence that auditorily induced semantic top-down information to early visual cortex actually impacts on visual perception. One reason for this lack of evidence might be that with conventional fMRI methods, top-down influences are difficult to detect when

early visual cortex is driven by feed-forward visual stimulation. Using more sensitive decoding fMRI methods, auditory top-down influences can be detected either in the absence of feedforward stimulation (Vetter et al., 2014) or, in the presence of visual stimulation, sharpen visual representations in early visual cortex (de Haas et al., 2013). When visual stimulation is ambiguous, semantically matching sound information can be discriminated in early visual cortex when it helps resolving the visual ambiguity (Vetter et al., 2016).

Let us summarize the challenge that the above reviewed evidence makes for both versions of the impenetrability claim: The weak impenetrability claim has received rather strong counterevidence from observations of (1) extensive structural feedback organization of the brain, (2) temporally very early feedback loops and (3) functional top-down processes modulating early visual processes by category-specific information. The strong impenetrability claim could incorporate these data by widening the “perceptual module” such that it includes rich but still internal processing in a very large perceptual module: Such a claim must or even wants to allow for “smart” perception (e.g. including category-specific information) but at the same time needs to demonstrate that all this smartness of perception is radically (and always) independent from any top-down influences originating in cognitive processing areas. This latter view presupposes an implausible version of a module.

3. Experimental evidence for cognitive penetration of visual experience

While the above evidence mainly speaks against the weak impenetrability claim and challenges the strong impenetrability claim, it still lacks the positive evidence to demonstrate cognitive penetration of actual perceptual experience. In the following section we focus on studies that clearly show a modification of perceptual color experience by activating abstract concepts and we aim to argue that these should best be explained as cases of cognitive penetration. As we will highlight, we may not be able to exclude the interpretation of perceptual learning at the moment, i.e. the claim of defenders of an impenetrability claim that all reported cases rely on long-term learning processes which actually modified the perceptual system and thus do not prove cognitive penetration. This analysis will motivate the discussion of short-term influences of higher cognitive processes in section 4.

3.1 Visual experience of colors can be affected by color concepts or by visual templates

The original idea of cognitively modified color experiences goes back to an old experiment by Delk and Fillenbaum (1965). Here, participants adjusted the background color for cut-out object shapes with typical or non-typical colors (red for a heart shape or color-neutral objects like a square or mushroom). While the cut-out objects always consisted of the same orange sheet of paper, subjects adjusted the color of the background as more red when the object was typically associated with a red color (e.g. a heart or lip) than when it was not. The study suggests that participants perceive the color of an object as more red if the shape of the object is associated with a red color. As such, this phenomenon can be seen as an example for cognitive penetration: the cognitive high-level knowledge that a heart shape is usually red changes basic color perception.

A modern version of this experiment with more rigorous methodology was carried out by Hansen et al. (2006). Here, participants were shown color photographs of fruits with a typical color, e.g. a yellow banana, on a grey background on a computer monitor. Participants adjusted the color of the fruit online until it appeared neutrally grey to them (while in another task they adjusted a grey banana until it appeared to have its natural color). Intriguingly, participants adjusted the banana not to its objectively achromatic grey, but to a slightly bluish grey shade – a grey opposite of yellow in color space. Thus, the objectively achromatic grey banana still appeared slightly yellow, and therefore participants adjusted it more towards blue so that it appeared neutrally grey. This effect was found not just for bananas but also for other typically colored fruits (e.g. orange, lemon, lettuce). These results nicely demonstrate that the color that is abstractly associated with an object shape affects very basic color perception, even if the object itself is shown in grey. We take this as a powerful example for cognitive penetration of perceptual experience: the abstract memorised knowledge of a banana typically being yellow affects the basic perceptual experience of color. The advantage of the study by Hansen et al. (2006) is that participants adjusted the color of the object online, while seeing the object without time constraints. Thus, the study measured direct and veridical perceptual experience without relying on memory or matching the percept with a reference.

Furthermore, this same type of experiment was done with the aim to investigate what we called the functional evidence for cognitive penetration: The influence of object identity on color perception even if objects are displayed achromatically also has a neural correlate all the way down in primary visual cortex V1. Bannert and Bartels (2013) showed in an elegant study that the yellow color of a grey banana is represented in the neural activity patterns of V1 even if all participants are seeing is a grey banana. That is, even if the feed-forward information to the visual system is achromatic grey, the typically associated color information of an object is communicated all the way down the visual hierarchy to V1. That is, abstract object identity affects veridical perceptual experience and has a neural consequence at the earliest neural processing level. Both speak in favour of cognitive penetration of perceptual experience and against an impenetrable early vision module.

How can defenders of cognitive impenetrability try to react to these observations? (a) They could argue that cognitive influences occurred at the level of judgment rather than at the level of perceptual experience because participants had to “judge” the color of an object when adjusting the color of object and background (e.g. in the study of Hansen et al. 2006). This argument is not convincing because what we deal with here are very basic online color comparisons. If one insists to think of them as perceptual judgments then it is important to consider that they are still done under normal perceptual conditions with unambiguous stimuli and without the need to learn new concepts. We know that basic evaluations of colors under these conditions are very reliable. Additional arguments would be needed to claim that under such conditions basic color judgements are systematically disturbed while perceptual experience remains unaffected (of course, with ambiguous stimuli or under suboptimal perceptual conditions it would be much more plausible to suppose cognitive penetration of judgement rather than perceptual experience).

(b) The alternative strategy to defend cognitive impenetrability could be to accept that this series of experiments demonstrates an actual influence on perceptual experience, but to deny it

as an argument against impenetrability because all influences take place within a single perceptual module and do not involve any top-down influences from cognitive areas. This argument only works for the strong impenetrability position. If one accepts these experiments as evidence for modifying perceptual experience, then the weak impenetrability claim cannot be maintained either because V1 is influenced by abstract concepts or visual templates (e.g. of yellow bananas), and these are too rich to be processed solely in early visual areas. The strong impenetrability claim in principle allows for top-down effects from rich abstract concepts onto “earlier” visual processes *as long as* they are completely part of the perceptual module. However, the latter presupposition is exactly what has never been established. After having presented striking evidence against the weak impenetrability claim (top-down influences on color experience and color processing), we can now focus on evidence against the strong impenetrability claim according to which all processing that happens before perceptual experience does not involve any cognition.

3.2 Visual experience of a scene is affected by activated memorized visual templates

Striking evidence for a change of perceptual experience is provided by impoverished black and white images, e.g. the well-known case of the Dalmatian dog. The image only seems to consist of randomly distributed black and white dots until we know to look out for a Dalmatian dog. Suddenly, we develop a different perceptual experience: we recognize the dog and can never go back to seeing the picture just as randomly distributed dots.

How can or should we interpret this change of perceptual experience while looking at the same visual scene under the same perceptual conditions? Since it seems that the only change from one situation to a second later is the activation of the concept *>Dalmatian dog<* that causes the dramatic change in the perceptual experience, the best interpretation seems to be that this is a case of cognitive penetration. That is, a (relevant) causal effect of higher cognitive contents, the concept of *>Dalmatian dog<*, probably involving a prototypical visual template of a Dalmatian dog, affects our perceptual experience. How is it possible to argue against cognitive penetration here? One way would be to argue that the activation of the concept and related visual template changes attentional processing and thereby produces new sensory inputs and thus influences perceptual experience. We do not want to deny that the activation of concepts may actually modify attentional processing but it needs to be shown by the defenders of an impenetrability claim that this is all that happens. Recent studies indicate that we can distinguish attention and expectation in visual experience and that top-down influences plausibly modify expectations playing a crucial role in perception (Summerfield & Egnér 2009). It is much more plausible that, in addition to changes in attentional processing, *processes of cognitive integration* of the black and white dots take place after early visual processes are completed but before perceptual experience is stable. The integration process leads from spatially ordered black and white dots to the image of a Dalmatian dog. This integration process corresponds to the recurrent feed-forward and feedback processing loops as already described in section 2.2. In addition to the empirical evidence reviewed there, an older study from Frith & Dolan (1997) provides nice support for the idea of cognitive integration of impoverished stimuli. They presented participants with an impoverished black and white image of a banana on some background.

Usually, hardly anyone recognizes the banana but perceives a pattern of black and white patches that cannot be integrated into any meaningful image. Later on, participants are presented with a clear image of the banana before viewing the impoverished image again. Contrasting the fMRI signal of the perception of the impoverished image before and after it was paired with the clear image resulted in a significant activation of the medial parietal lobe. The medial parietal lobe cannot be regarded as a candidate for early visual processes; it is thus most plausibly a candidate for higher-level processing. Therefore, changes of the perceptual experience of impoverished images are best explained as a result of cognitive penetration from a high-level area. The only alternative would be to include the medial parietal lobe into a big perceptual module which again is an implausible claim as discussed above.

There is a final important defense strategy used by the impenetrability proponents, i.e. the perceptual learning argument. If they concede a modification of the perceptual processes happening between sensory input and perceptual experience, then this is supposedly a result of long-term changes within the perceptual module due to associative learning; e.g. the perceptual system supposedly learns what a Dalmatian dog or a banana looks like independent of any input from higher cognitive processes. In our view, we need at least the situational activation of the template of the Dalmatian dog to account for the observed change of perceptual experience. How should we classify this template activation: can it be purely perceptual or does it involve a higher cognitive process? The visual template of a Dalmatian dog is a rich and categorical template with quite some variation (prototype character). This prototypical or invariant representation is plausibly anchored in higher cognitive processes.

Furthermore, visual templates of objects (such as the prototypical template of a Dalmatian dog) are unlikely to be created solely by early visual processes: at least in the case of children (and adults, of course) who learned the concept of a Dalmatian dog, the visual template is part of this concept and intensely interwoven with higher cognitive processes constituting the concept (this is a consequence of the discovery that concepts are anchored in sensorimotor representations, Barsalou 1999, 2008). An isolated activation of a parsimonious visual template which would be a candidate for perceptual learning is very implausible because concepts involve a rich organizational structure (see also Newen & Marchi, 2016). For the weak impenetrability claim this would exclude a pure perceptual learning explanation. Then, we are left with the strong impenetrability claim and thus, we are back to the arguments on the implausibility of a big perceptual module (see above) encompassing creation, storage and retrieval of visual templates and abstract concepts such as “animal”, “dog”, “fruit” or “banana” (and all other possible object categories). Additional convincing evidence against a perceptual learning explanation could be provided if there was evidence that perceptual experience can not only be modified by concepts acquired through long-term learning but by concepts implemented on a *short-term time scale*. For example, Kok et al. (2012) associated tones to grating stimuli of different orientations, and could thus elicit a visual expectation or activation of a specific visual template upon playing the tone to participants. The expected stimulus resulted in higher orientation sensitivity and sharpened neural representations in V1. Thus, priming a visual stimulus with an acoustic tone leads to template activation (expected orientation of the grating) and this modifies the perceptual processes as early as V1. One may still argue that these short-term modifications are part of a limited visual module because only

simple gratings (of orientation and contrast) were used which may not necessarily involve higher-level cognitive processes or concepts. Therefore we report in the following a study dealing with the recognition of numbers which is clearly a conceptual process. This fascinating experiment used hypnosis. Hypnotic suggestions are useful to implement new abstract concepts or visual templates on a short time scale. These concepts or templates can modify perceptual experience and their implementation. The activation of these concepts cannot be explained by perceptual learning since this is supposed to be a long-term process.

3.3. Visual experience can be affected by posthypnotic suggestions

With the aim to investigate the phenomenon of synesthesia, Cohen Kadosh et al. (2009) used hypnosis to test whether experiences of synesthesia could be induced in non-synesthetes. Although not being the focus of the study, the results indicate that cognitive penetration of perceptual experience by newly acquired cognitive contents, excluding perceptual learning, is indeed possible. One group of participants was hypnotized with the aim of inducing digit-color associations such as “1” always being red, “2” always being yellow. Then participants were taken out of hypnosis (not consciously remembering the digit-color associations) and were told the following posthypnotic suggestion: “Look at that color; this is the color of the digit X, and whenever you see, think, or imagine that digit, you will always perceive it in that color.” After this suggestion, they were requested to perform a simple perceptual detection task of a digit printed in black. Let us explain the experiment in the case of the suggestion that “2” always appears yellow. The display actually presented a black “2” on differently colored backgrounds excluding black (red, yellow, green, blue, etc.). The working hypothesis was that participants with posthypnotic suggestion should not perceive the “2” on a yellow background because, according to the suggestion, the “2” should have the same color. The results of the study confirmed this hypothesis: participants who received posthypnotic suggestions of specific digit-color association did not perceive the black digit associated with the same color as the background. Without post-hypnotic suggestion, the same participants perceived the digit normally. One conclusion of this study is that the semantic content of a short-term and reversible posthypnotic suggestion caused a change in perceptual experience. This rules out long-term perceptual learning effects as a potential cause of the observed effect. Furthermore, this study demonstrates a change of actual perceptual experience (rather than judgement) because the effect is based on a simple perceptual detection task.

3.4 Intermediate summary

We presented evidence for each of the key features indicating cognitive penetration: 1. We have shown that recent discoveries from neuroscience support top-down influences on early visual processes (including V1), including top-down influences with categorical content. We argued from three different perspectives: (a) brain structure and organisation, (b) temporal processing and (c) the functional evidence for categorical top-down influences. Since not all this evidence is directly related to visual experience, we reported in section 3 the most challenging studies demonstrating modified visual experiences by activating a cognitive content. The studies on modified color experiences provide evidence against the weak impenetrability claim. The

observations on impoverished black and white stimuli (sec. 3.2) provide evidence against the strong impenetrability thesis: cognitively triggered integration processes can play an important role in producing the perceptual experience of a visual scene. Finally, we excluded the argument from perceptual learning, i.e. that all reported effects do not involve cognitive top-down processes because they are supposed to be only a consequence of long-term learning processes that modified the perceptual system. This interpretation is excluded by the short-term cognitive effects shown in the hypnosis study (sec. 3.3). Thus, the sum of evidences justifies the claim that cognitive penetrability is the most plausible claim based on an inference to the best explanation (or systematization). At least the burden of proof lies now clearly on the side of the impenetrability claim.

Defenders of the impenetrability claim could still argue that all this empirical evidence is not provided by a single experiment but by many different ones. This is not really a criticism as in empirical science, it is mostly the critical mass of evidence from many different studies complementing and replicating each other that form scientific facts. However, one could still wish for a decisive knock-down experiment. Ideally, such an experiment combines (i) a clear demonstration of a change of perceptual experience with a perceptual detection task, (ii) structural and functional evidence from neurosciences supporting the modification of perceptual processing, (iii) a demonstration that the effects are produced by a higher cognitive process, and finally (iv) that this cognitive process is implemented on a short time scale and thus cannot be explained away as an effect of perceptual learning. This is the methodological step to be taken by the ideal experiment demonstrating cognitive penetration. The evidence already available indicates that we are on our way to get there.

4. Perception and Cognition: Is there a boundary between them?

By defending cognitive penetration, do we deny any boundary between perception and cognition? Although this may seem so at first glance, this is not the case. To the contrary, without distinguishing between perception and cognition, we cannot get the debate started at all. We have to presuppose at least some distinction to ask the question whether higher-level *cognitive* processes can causally influence *perceptual* processes leading to a perceptual experience. Thus, to deny cognitive penetration because of the danger of losing a distinction between perception and cognition is ill motivated. We can keep some distinction but we have to give up the idea of a clear-cut distinction.

How should we think of the relation between cognition and perception then? Before specifying our own view, let us shortly report the most extreme positions that have recently been discussed. Firestone and Scholl (2015) combine their denial of cognitive penetration with the classic position that there is a sharp boundary between perception and cognition such that cognition never affects perception: “We have argued that there is a joint between perception and cognition to be carved by cognitive science, and that the nature of this joint is such that perception proceeds without any direct, unmediated influence from cognition.” As we have argued, this radical view of a large, totally non-penetrable perceptual module is neither compatible with empirical evidence nor is it conceptually or empirically plausible to speak of such a large perceptual module. The extreme counter position is e.g. expressed by Andy Clark (2013):

“All this makes the lines between perception and cognition fuzzy, perhaps even vanishing. In place of any real distinction between perception and belief we now get variable differences in the mixture of top-down and bottom-up influence, and differences of temporal and spatial scale in the internal models that are making the predictions. Top level (more ‘cognitive’) models intuitively correspond to increasingly abstract conceptions of the world and these tend to capture or depend upon regularities at larger temporal and spatial scales. Lower level (more ‘perceptual’) ones capture or depend upon the kinds of scale and detail most strongly associated with specific kinds of perceptual contact. (...) To perceive the world just is to use what you know to explain away the sensory signal across multiple spatial and temporal scales. The process of perception is thus inseparable from rational (broadly Bayesian) processes of belief fixation, and context (top down) effects are felt at every intermediate level of processing. As thought, sensing, and movement here unfold, we discover no stable or well-specified interface or interfaces between cognition and perception.” (2013: 10)

In Clark’s view, we should give up any distinction between perception and cognition because the processing of the brain is best modelled as one of predictive coding all the way up and all the way down. In the framework of predictive coding it seems to make no sense to introduce any boundary. If one takes this position seriously, it follows that we no longer should distinguish between perception and cognition at all. But this would not allow us to account for visual illusions which remain stable even if we have full knowledge about how the perceptual system is tricked into the illusion. We highlighted that these phenomena are not the standard cases of everyday perception but they clearly exist. It seems to be that Clark is throwing the baby out with the bathwater. How can we account for stable visual illusions, on the one hand, and for cognitive penetration, on the other, and still tell a plausible story of the relation between perception and cognition?

We agree that predictive coding (Clark 2013; Hohwy 2014) is a very plausible framework for explaining brain processing and perception. It implies a hierarchy of processes without any sharp boundary between these processes. But this framework is very general and remains compatible with distinguishing paradigmatic phenomena of cognition and perception even if the majority of phenomena lies in the middle ground in between both. The need to distinguish these three cases is exactly the account we take to be adequate: pure perception, everyday perception and pure cognition (including perception-based judgment). Paradigmatic phenomena of cognition are linguistic judgments but also perception-based beliefs or judgments which are conceptual but not necessarily linguistic (Newen & Bartels 2007): the content of a judgment can be completely characterized by the composition of conceptual contents. Paradigmatic cases of (pure) perception are e.g. the perception of impoverished pictures of black and white dots: the content of the perception of such pictures can best be described as non-conceptual, spatial organization of these dots independent from any high-level conceptual contents. The middle ground includes the majority of everyday perceptual experiences: to describe the content of these experiences, it seems adequate to use a combination of descriptions of non-conceptual, spatial organizations of dots intertwined with high-level and abstract conceptual contents to create a meaningful percept. Everyday perceptual experiences can have rich contents (Newen, 2016) but nevertheless can be clearly distinguished from perception-based judgments. Our suggested account of distinguishing paradigmatic pure perception from paradigmatic pure cognition with the majority of perceptual phenomena as everyday rich perception as a middle ground presupposes the possibility of rich perceptual contents (Newen,

2016) and allows for both: that these rich contents are produced by perceptual learning or by cognitive penetration.

5. Conclusion

Let us summarize our view: We presented striking evidence against the weak impenetrability claim including (1) extensive structural feedback organization of the brain, (2) temporally very early feedback loops and (3) functional top-down processes modulating early visual processes by category-specific information. The strong impenetrability claim could incorporate these data by widening the “perceptual module” such that it includes rich but still internal processing in a very large perceptual module: Such a claim must or even wants to allow for “smart” perception (e.g. including category-specific information) but at the same time needs to demonstrate that all this smartness of perception is radically (and always) independent from any top-down influences originating in cognitive processing areas. We argued that this latter view presupposes an implausible version of a module. Therefore, we have to accept cognitive penetration of our perceptual experience. Furthermore, we described the framework for an ideal experiment demonstrating cognitive penetration as a future knock-down argument in the debate. From our perspective the evidence from many independent studies indicate that this will be manageable in the near future. Finally we outlined that our position against cognitive impenetrability does not imply that we have to give up the distinction between perception and cognition completely. A plausible view nicely compatible with cognitive penetration distinguishes paradigmatic cases (i) of perception, (ii) of cognition and allows for (iii) intermediate phenomena with a strong intertwinement of perception and cognition with everyday perception being the strongest candidate.

Acknowledgement

P.V. was supported by a fellowship from the German Research Foundation (Ve739/1-1).

References

- Bannert, M. M., & Bartels, A. (2013). Decoding the yellow of a gray banana. *Current Biology*, 23(22), 2268–2272.
- Barsalou, L.W. (1999). Perceptual symbol system. *Behavioral and Brain Science* 22(4), 637–660.
- Barsalou, L.W. (2008). Cognitive and neural contributions to understating the conceptual system. *Current Directions in Psychological Science* 17(2), 91-95.
- Ber, A. L., Plank, T., & Greenlee, M. W. (2011). Diffusion tensor imaging shows white matter tracts between human auditory and visual cortex. *Experimental Brain Research*, 213(2–3), 299–308.
- Carandini, M., Demb, J. B., Mante, V., Tolhurst, D. J., Dan, Y., Olshausen, B. A., Gallant, J. L., & Rust, N. C. (2005). Do we know what the early visual system does? *Journal of Neuroscience*, 25(46), 10577–10597.

- Chen, Y.-C., & Spence, C. (2011a). Crossmodal semantic priming by naturalistic sounds and spoken words enhances visual sensitivity. *Journal of Experimental Psychology. Human Perception and Performance*, 37(5), 1554–1568.
- Chen, Y.-C., & Spence, C. (2011b). The crossmodal facilitation of visual object representations by sound: evidence from the backward masking paradigm. *Journal of Experimental Psychology. Human Perception and Performance*, 37(6), 1784–1802.
- Chen, Y.-C., Yeh, S.-L., & Spence, C. (2011). Crossmodal constraints on human perceptual awareness: auditory semantic modulation of binocular rivalry. *Frontiers in Psychology*, 2, 212.
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3), 181–204.
- Cohen Kadosh, R., Henik, A., Catena, A., Walsh, V., & Fuentes, L. J. (2009). Induced cross-modal synaesthetic experience without abnormal neuronal connections. *Psychological Science*, 20(2), 258–265.
- Çukur, T., Huth, A. G., Nishimoto, S., & Gallant, J. L. (2013). Functional subdomains within human FFA. *Journal of Neuroscience*, 33(42), 16748–16766.
- David, S. V., & Gallant, J. L. (2005). Predicting neuronal responses during natural vision. *Network*, 16(2–3), 239–260.
- de Haas, B., Schwarzkopf, D. S., Uner, M., & Rees, G. (2013). Auditory modulation of visual stimulus encoding in human retinotopic cortex. *NeuroImage*, 70, 258–267.
- Delk, J. L., & Fillenbaum, S. (1965). Differences in perceived color as a function of characteristic color. *The American Journal of Psychology*, 78(2), 290–293.
- Dretske, F. (1981). *Knowledge and the flow of information*. Cambridge/London: MIT Press.
- Drewes, J., Goren, G., Zhu, W., & Elder, J. H. (2016). Recurrent processing in the formation of shape percepts. *Journal of Neuroscience*, 36(1), 185–192.
- Eckert, M. A., Kamdar, N. V., Chang, C. E., Beckmann, C. F., Greicius, M. D., & Menon, V. (2008). A cross-modal system linking primary auditory and visual cortices: Evidence from intrinsic fMRI connectivity analysis. *Human Brain Mapping*, 29(7), 848–857.
- Firestone, C., & Scholl, B. J. (2015). Cognition does not affect perception: Evaluating the evidence for “top-down” effects. *Behavioral and Brain Sciences*, 1–77.
- Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.
- Foxe, J. J., & Simpson, G. V. (2002). Flow of activation from V1 to frontal cortex in humans. A framework for defining “early” visual processing. *Experimental Brain Research*, 142(1), 139–150.
- Frith, C., & Dolan, R. J. (1997). Brain mechanisms associated with top-down processes in perception. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 352(1358), 1221–1230.
- Gauthier, I., Skudlarski, P., Gore, J. C., & Anderson, A. W. (2000). Expertise for cars and birds recruits brain areas involved in face recognition. *Nature Neuroscience*, 3(2), 191–197.
- Gilbert, C. D., & Li, W. (2013). Top-down influences on visual processing. *Nature Reviews. Neuroscience*, 14(5), 350–363.
- Hansen, T., Olkkonen, M., Walter, S., & Gegenfurtner, K. R. (2006). Memory modulates color appearance. *Nature Neuroscience*, 9(11), 1367–1368.

- Howhy, J. (2014). *The predictive mind*. Oxford: Oxford University Press.
- Hupé, J. M., James, A. C., Girard, P., Lomber, S. G., Payne, B. R., & Bullier, J. (2001). Feedback connections act on the early part of the responses in monkey visual cortex. *Journal of Neurophysiology*, *85*(1), 134–145.
- Hupé, J. M., James, A. C., Payne, B. R., Lomber, S. G., Girard, P., & Bullier, J. (1998). Cortical feedback improves discrimination between figure and background by V1, V2 and V3 neurons. *Nature*, *394*(6695), 784–787.
- Huth, A. G., Nishimoto, S., Vu, A. T., & Gallant, J. L. (2012). A continuous semantic space describes the representation of thousands of object and action categories across the human brain. *Neuron*, *76*(6), 1210–1224.
- Kanwisher, N., Stanley, D., & Harris, A. (1999). The fusiform face area is selective for faces not animals. *Neuroreport*, *10*(1), 183–187.
- Kanwisher, N., & Yovel, G. (2006). The fusiform face area: a cortical region specialized for the perception of faces. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *361*(1476), 2109–2128.
- Kok, P., Jehee, J. F. M., & de Lange, F. P. (2012). Less is more: expectation sharpens representations in the primary visual cortex. *Neuron*, *75*(2), 265–270.
- Lamme, V. A., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences*, *23*(11), 571–579.
- Macpherson, F. (2012). Cognitive penetration of color experience. Rethinking the issue in light of an indirect mechanism. *Philosophy and Phenomenological Research*, *84*(1), 24–62.
- Markov, N. T., Ercsey-Ravasz, M., Van Essen, D. C., Knoblauch, K., Toroczkai, Z., & Kennedy, H. (2013). Cortical high-density counterstream architectures. *Science*, *342*(6158), 1238406.
- Michel, C. M., Seeck, M., & Murray, M. M. (2004). The speed of visual cognition. *Supplements to Clinical Neurophysiology*, *57*, 617–627.
- Morishima, Y., Akaishi, R., Yamada, Y., Okuda, J., Toma, K., & Sakai, K. (2009). Task-specific signal transmission from prefrontal cortex in visual selective attention. *Nature Neuroscience*, *12*(1), 85–91.
- Newen, A. (2016). Defending the liberal-content view of perceptual experience: direct social perception of emotions and person impressions. *Synthese*, doi: 10.1007/s11229-016-1030-1033.
- Newen, A. & Bartels, A. (2007). Animal minds and the possession of concepts. In: *Philosophical Psychology* *20*, 283–308.
- Newen, A. & Marchi, F. (2016). Concepts and their organizational structure: Concepts are templates based on mental files. In D. Hommen, C. Kann, T. Osswald (Eds), *Concepts and Categorization*. (pp. 197–227). Münster: mentis Verlag.
- O’Shea, J., Muggleton, N. G., Cowey, A., & Walsh, V. (2004). Timing of target discrimination in human frontal eye fields. *Journal of Cognitive Neuroscience*, *16*(6), 1060–1067.
- Pascual-Leone, A., & Walsh, V. (2001). Fast backprojections from the motion to the primary visual area necessary for visual awareness. *Science*, *292*(5516), 510–512.

- Pearson, J., Naselaris, T., Holmes, E. A., & Kosslyn, S. M. (2015). Mental Imagery: Functional Mechanisms and Clinical Applications. *Trends in Cognitive Sciences*, 19(10), 590–602.
- Petro, L. S., Vizioli, L., & Muckli, L. (2014). Contributions of cortical feedback to sensory processing in primary visual cortex. *Frontiers in Psychology*, 5, 1223.
- Plomp, G., Hervais-Adelman, A., Astolfi, L., & Michel, C. M. (2015). Early recurrence and ongoing parietal driving during elementary visual processing. *Scientific Reports*, 5, 18733.
- Pylyshyn, Z. W. (1999). Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, 22, 341–365.
- Raftopoulos, A. (2011). Late vision: processes and epistemic status. *Frontiers in Psychology*. doi.org/10.3389/fpsyg.2011.00382.
- Raftopoulos, A. (2014). The cognitive impenetrability of the content of early vision is a necessary and sufficient condition for purely nonconceptual content. *Philosophical Psychology*, 27(5), 601–620.
- Schirmer, A., Fox, P. M., & Grandjean, D. (2012). On the spatial organization of sound processing in the human temporal lobe: A meta-analysis. *NeuroImage*, 63(1), 137–147.
- Scholte, H. S., Jolij, J., Fahrenfort, J. J., & Lamme, V. A. F. (2008). Feedforward and recurrent processing in scene segmentation: Electroencephalography and functional magnetic resonance imaging. *Journal of Cognitive Neuroscience*, 20(11), 2097–2109.
- Siegel, S. (2005). Which properties are represented in perception? In T. Szabó Gendler, & J. Hawthorne (Eds.), *Perceptual experience* (pp. 481–503). Oxford: Clarendon.
- Silvanto, J., Cowey, A., Lavie, N., & Walsh, V. (2005). Striate cortex (V1) activity gates awareness of motion. *Nature Neuroscience*, 8(2), 143–144.
- Silvanto, J., Lavie, N., & Walsh, V. (2006). Stimulation of the human frontal eye fields modulates sensitivity of extrastriate visual cortex. *Journal of Neurophysiology*, 96(2), 941–945.
- Summerfield, C., & Egnér, T. (2009). Expectation (and attention) in visual cognition. *Trends in Cognitive Sciences*, 13(9), 403–409.
- Vetter, P., Grosbras, M.-H., & Muckli, L. (2015). TMS over V5 disrupts motion prediction. *Cerebral Cortex*, 25(4), 1052–1059.
- Vetter, P., Petrini, K., Piwek, L., Smith, F.W., Solanki, V., Bennett, M., Pollick, F. & Muckli, L. (2016) Decoding emotional valence of sounds in early visual cortex. *J Vis (Annual Meeting of the Vision Sciences Society)*.
- Vetter, P., Smith, F. W., & Muckli, L. (2014). Decoding sound and imagery content in early visual cortex. *Current Biology*, 24(11), 1256–1262.