How to (and how not to) think about top-down influences on visual perception

Christoph Teufel[1] & Bence Nanay[2]

[1] School of Psychology and Cardiff University Brain Research Imaging Centre, Cardiff University, 70 Park Place, CF10 3AT, Wales, UK

[1] Centre for Philosophical Psychology, University of Antwerp, D 413, Grote Kauwenberg 18, 2000 Antwerp, Belgium

Abstract
The question of whether cognition can influence perception has a long history in neuroscience and philosophy. Here, we outline a novel approach to this issue, arguing that it should be viewed within the framework of top-down information-processing. This approach leads to a reversal of the standard explanatory order of the cognitive penetration debate: we suggest studying top-down processing at various levels without preconceptions of perception or cognition. Once a clear picture has emerged about which processes have influences on those at lower levels, we can re-address the extent to which they should be considered perceptual or cognitive. Using top-down processing *within* the visual system as a model for higher-level influences, we argue that the current evidence indicates clear constraints on top-down influences at all stages of information processing; it does, however, not support the notion of a boundary between specific types of information-processing as proposed by the cognitive impenetrability hypothesis.

One influential debate about perception concerns its purity: is perception an encapsulated process that is protected from influences by cognition or is perceptual bottom-up processing influenced by top-down cognitive information? This debate, which is frequently referred to as the 'cognitive penetration debate', is complicated by the fact that it is often not clear what kind of mental state is supposed to be doing the penetrating and what kind of mental state is supposed to be penetrated. In other words, it is not clear what is 'top' and what is 'bottom' in the debate about top-down influences on perception.

In the first part of this paper, we attempt to clarify some of these conceptual issues. We then proceed to suggest a practical alternative to the philosophical turf wars concerning the extent to which perception is encapsulated. We start with two uncontroversial observations: First, even the most avid proponents of the view that perception is cognitively impenetrable accept the existence of top-down processing *within* the visual system;[1] in other words, it is uncontroversial that higher-levels of visual processing feed back information to and shape lower-levels of visual processing (Pylyshyn, 1999). And second, the exact locus of the boundary between perception and cognition is notoriously difficult to determine (Masrour 2011, Siewert 2002, Siegel 2007, Kriegel 2007 and Bayne 2009, Nanay 2011, 2012, 2013). On the basis of these two observations, we argue that important insights might be gained once we stop focusing exclusively on top-down modulation of perception by cognition; rather, we suggest that it is heuristically valuable to view this special case within the broader context of top-down influences in a hierarchically organised information processing system.

The general agreement concerning top-down processing within the visual system can be used as a starting point from which our understanding can be expanded to potential higher-level top-down influences without having to commit *a priori* to what exactly counts as perception or cognition. Once such an approach has been adopted, we can start asking nuanced questions about the specific mechanism of top-down modulation in information processing in general. Among other questions, we can ask what kinds of constraints exist on top-down influences between certain levels of processing and whether some of these constraints might amount to a full-blown boundary as proposed by the encapsulation hypothesis. According to our evaluation of the theoretical and empirical evidence, there is no reason to assume that top-down processing is restricted to specific parts of the information processing hierarchy. By contrast, we defend a view that puts clear constraints on all sorts of top-down processing – as well as on bottom-up processing – but that allows bidirectional flow of information between levels that some would consider to cross the perception-cognition divide.

I. Two debates about top-down influences on perception

The main conceptual confusion concerning debates about top-down influences on perception is that it is not clear what is meant by 'perception' in this context. Many philosophers (Siegel 2011, Macpherson 2012, Stokes 2012), but also some

---

[1] We will focus on the visual sense modality in this paper, but there is no prima facie reason why our conclusions could not be generalized to other sense modalities.

psychologists (e.g., Firestone and Scholl 2016), take 'perception' to be perceptual experience: something we are consciously aware of. According to this conceptualisation, the question is whether top-down influences can alter the way we experience a scene – the phenomenal character of our experience: what it is like to perceive this scene.

Another way of understanding what is meant by 'perception' when we talk about top-down influences is perceptual processing – something neuroscientists, psychophysicists, and some psychologists worry about. Here, the question is whether a certain type of information processing is influenced in a top-down manner. Most generally, all computations that are specialised for and concerned with transforming the spatio-temporal pattern of light hitting the retina into meaningful representations can be considered part of visual processing. Following the work of David Marr (1982), visual processing is often specified in further detail in terms of computations that lead to specific geometrical descriptions of a visual scene (e.g. Pylyshyn, 1999). Our understanding of the neurobiological architecture of the early parts of the visual system is advanced enough to provide a general picture of how such a functional description of vision is realised by populations of neurons that extract information about specific perceptual properties such as orientation, contours, motion, etc. in striate and extrastriate cortices. Evidence of top-down modulation in these brain areas would therefore count as evidence for top-down modulation of perceptual processing.

These two questions are clearly very different – one of them is about phenomenology and the other is about a specific type of information processing. We are extremely pessimistic about whether the first of these debates could ever be resolved in a satisfactory manner. The main reason for this pessimism is methodological in nature. Most studies that attempt to address the extent to which perceptual phenomenology is influenced by top-down processing rely, at least to some degree, on introspection, which is known to be notoriously unreliable (see Nisbett and Wilson 1977, Kahneman and Tversky 1973, Tversky and Kahneman 1981, Zhong and Liljenquist 2006, Williams and Bargh 2008, Wegner 2002, Haggard et al. 2002, Wegner et al. 2004, Greenwald and Banaji 1995, Dunham et al. 2008 and Schwitzgebel 2008, Spener and Bayne 2010 – this is merely the tip of the iceberg of the vast literature on the unreliability of introspection). But we do not mean to suggest that it is *only* by introspection that one can find out about perceptual phenomenology. The so-called 'methodology of contrast cases' (Siegel 2006, Kriegel 2007), for example, combines introspective evidence with an inference to the best explanation – it is not introspection alone that does the job (although introspection is a necessary ingredient of all attempts to characterize perceptual phenomenology). What we take to be an even more important worry about the focus on perceptual phenomenology is that it is difficult to settle disagreements about the nature of perceptual phenomenology.

The main difficulty in this debate is to determine what is part of our perceptual as opposed to non-perceptual phenomenology. Those who argue for the existence of top-down influences on perceptual phenomenology need to show

that there can be two mental states, call them M1 and M2, that only differ in that there is a top-down influence in M2, which is missing in M1 and that the two differ in their perceptual phenomenology. So the top-down influence results in a difference in perceptual phenomenology. Those who are against the idea of top-down influences on perceptual phenomenology can acknowledge that M1 and M2 differ only in that top-down influences are present in M2 but absent in M1 and they can also acknowledge that M1 and M2 differ in their non-perceptual phenomenology – they only deny that they differ in their perceptual phenomenology. So the only way of adjudicating between the proponents and the opponents of top-down influences on perceptual experience is by having a very clear distinction between perceptual and non-perceptual phenomenology.

But we are blatantly missing any such clear distinction or even a methodology that might be able to help develop it. Take the following example: You are at a dinner party and are eating what you take to be chicken. Then your host tells you that it is in fact rat meat. Your experience, presumably, changes. The meat tastes differently. This seems to be an indication that your perceptual phenomenology changes – what changes is the way the meat tastes to you. But suppose that you insist that what changed was not the perceptual but the non-perceptual phenomenology in this example. It is difficult to see what could possibly settle this disagreement. We may be able to tell whether our overall phenomenology changed. But to tell whether this phenomenal change was perceptual or non-perceptual is much more difficult. In other words, if I say that M1 and M2 differ in their perceptual phenomenology and you deny this, it is not clear how the issue can be decided or what methodology might provide further insight. Intuitions wildly differ with regards to what phenomenal character counts as perceptual.

The debate about whether there are top-down influences on perceptual experiences inherits the problems that arise from the lack of a methodology to tell perceptual and non-perceptual phenomenology apart. However, once we accept a conceptualisation of perception in terms of perceptual processing, the situation substantially improves. Here, we can draw on a rich methodological toolbox that includes psychophysical techniques, electrophysiology, neuroimaging, and computational neuroscience. While each of these methodologies comes with its own limitations and difficulties, they generally provide sophisticated ways of shedding light on perceptual processing at different levels of description from single neurons to full perceptual representations.

As a supplementary point, we would like to briefly highlight one specific discussion within psychophysics, which has not been explored in much detail by philosophers but, we believe, could have interesting implications for the relationship between the two debates highlighted above ('perception as phenomenology' vs. 'perception as processing'). While psychophysicists are interested in perceptual processing, there is a certain distinction within psychophysics that, somewhat indirectly, alludes to phenomenology: the distinction between performance-based and appearance-based tasks (Kingdom & Prins, 2009), which is closely related to the distinction between Class A and

Class B observations (Brindley, 1960; Morgan, Melmoth, & Solomon, 2013). Without going into any of the technical details, roughly, performance-based tasks attempt to measure how good an observer is at a specific task, they measure the limits of sensitivity of perceptual processing. These aspects of perception can be characterised without any reference to phenomenology in human observers, and any biological or artificial measurement system. A typical example would be the characterisation of an observer's absolute contrast detection threshold, i.e., the lowest contrast an observer is able to detect at a certain performance level.

Appearance-based tasks on the other hand attempt to measure certain biases in perceptual processing. For instance, in the Müller-Lyer illusion the orientation of the fins with respect to the central line induce a certain bias in size processing and change how observers judge the length of the central line. It would not make sense to ask how sensitive the observer is to differences in length in this instance; rather, appearance-based tasks measure changes in the *apparent* magnitude of a certain stimulus dimension. While psychophysicists typically avoid specifying what exactly they mean by appearance or apparent magnitude of stimulus dimension, these tasks seem to focus on something similar to what philosophers have in mind when talking about perceptual phenomenology. Consequently, it is not surprising that psychophysical work that relies on appearance-based tasks has been haunted by problems that share some similarity to the ones mentioned above in the context of perceptual phenomenology. In particular, it is difficult to ensure that measured biases are perceptual in nature (Jogan & Stocker, 2014; Morgan et al., 2013; Morgan, Dillenburger, Raphael, & Solomon, 2012; Storrs, 2015).

In psychophysics, this problem is indirectly addressed by attempts to exclude the possibility for non-perceptual factors such as response biases and decision biases to affect results. In particular, in recent years, researchers have developed sophisticated methodologies to precisely control for the contamination of findings by such non-perceptual biases (Jogan & Stocker, 2014; Morgan et al., 2013). To our knowledge, these methods have so far not been applied to the study of top-down influences on perception; yet, given that many supposed top-down effects in perception that philosophers have been interested in can essentially be re-conceptualised as biases, we believe that these novel tasks will be of huge significance for new insights in this domain. More generally, we believe that closer analyses of the distinction between performance- and appearance-based tasks or Class A/B observations in psychophysics might hold some insights for philosophers interested in perception as phenomenology in the context of top-down effects.

To summarise this section, we believe that some of the conceptual confusion concerning top-down influences on visual perception arise from at least two distinct notions of what perception is considered to be in this context: perceptual phenomenology or perceptual information processing. Given the fundamental problems in distinguishing perceptual from non-perceptual phenomenology, we argue that the real debate about top-down influences on perception is about whether perceptual processing is influenced in a top-down manner. This is what

we take to be the question of 'top-down influences on perception' in what follows.[2]

II. Cognitive penetration and top-down processing

We view the question of cognitive penetration as part of the wider research endeavour that attempts to understand top-down influences on perception. We believe that this perspective has various advantages over an exclusive focus on cognitive influences on perceptual processing. The most important benefit of such an approach is that it provides us with a framework that allows us to ask more nuanced questions: the question of top-down influences on perception is no longer a simple yes-no question as in the case of cognitive penetration, but a multifaceted one. For instance, we can ask specific questions about which types of computations at which level of processing have effects on mechanisms further down the hierarchy. We can ask whether and what kinds of constraints are exerted onto top-down influences at different levels of processing. We can ask functional questions of why such constraints might or might not exist.

Moreover, when we start addressing these questions, we can build on a solid body of literature concerned with top-down effects: those *within* the visual system. We will provide a few examples of such effects in the next section. For the purpose of this section, it suffices to say that the existence of top-down influences from higher onto lower levels of processing *within* the visual system is accepted even by researchers, who strongly oppose the idea that vision can be influenced by cognition (Pylyshyn, 1999). These findings provide a natural starting point for the attempt to push the boundary of top-down effects up the hierarchy.

One benefit that follows from viewing cognitive penetration in the wider context of top-down influences on perception is that we can reverse the standard explanatory order of the cognitive penetration debate. The original way of raising the question of cognitive penetration was this: First, we need independent ways of identifying cognitive processing and perceptual processing. Second, we can ask about the influences between the two (namely, whether the former influences the latter). But we have seen that these independent ways of characterising perception and cognition are often based on unsatisfactory methodology. There are major differences in what different researchers consider to be perception or cognition, and some authors question the validity of the perception/cognition distinction altogether. In fact, within the cognitive penetration debate, it has been repeatedly suggested in the past couple of years that the abundance of top-down influences in perceptual and non-perceptual processing forces us to take perception to be continuous with cognition in the

---

[2] It is important to highlight that recent work by Firestone and Scholl (forthcoming) against top-down influences on perception concern very explicitly the debate about whether there are top-down influences on perceptual phenomenology. So their attack on top-down influences are strictly speaking irrelevant for our argument and they (as well as the methodological controversies noted in the commentaries to this article) provide an additional reason to shift the debate from perceptual phenomenology to perceptual processing.

sense that there is no real boundary between them (Clark, 2013; Fletcher & Frith, 2009, Lupyan 2010, 2015, Vetter and Newen 2014, Hohwy 2013).

It is important to point out that this somewhat radical conclusion does not follow from our proposal here. What we propose is a reverse order of explanation: without assuming anything about what is perception and what is cognition, we can identify top-down influences at a wide variety of points of information processing, from very low levels to very high levels (for an overview of some relevant studies see the following sections). And once we have a clearer picture of what types of computational mechanisms have influences on those at lower levels, we can re-address the question of the extent to which they should be considered perceptual or cognitive. If it turns out that there is a continuous cascade of top-down flow of information – and, in our reading, the current state of research seems to suggest exactly that – then no matter how we define perception and cognition (and regardless of where we draw the perception/cognition boundary), we have good reason to conclude that perception is subject to top-down cognitive influences. Note, however, that nothing in this proposal forces us to deny that there is a point in neural processing where strictly perceptual processing ends and cognitive processing begins.

No-one thinks that the fact that perception influences cognition implies that there is no difference between the two. Similarly, we can hold on to a distinction between perception and cognition even if there are influences from the latter onto the former. In other words, the question of whether there is a boundary between perception and cognition is orthogonal to the question of what kind of uni- or bidirectional influences exist between the two. To sum up, our way of conceptualising cognitive penetration as part of a continuous cascade of top-down processes that stretches from very high-level processing to very low-level processing is consistent with denying that there is any meaningful perception vs. cognition boundary. But it is also consistent with the existence of such a boundary.

III. Top-down processing within the visual system

In the past decades, our understanding of the structural and functional organisation of visual systems in various mammalian species has improved dramatically and, in nonhuman and human primates, an increasing number of cortical areas dedicated to processing visual inputs has been and still is being characterised (for reviews see (Bullier, 2004; Felleman & Van Essen, 1991; Grill-Spector & Malach, 2004; Katzner & Weigelt, 2013; Van Essen, 2004; Van Essen & Maunsell, 1983). In humans and nonhuman primates, the main visual pathway connects neural networks in the retina to the primary visual cortex (V1) via the lateral geniculate nucleus (LGN) in the thalamus; outputs from V1 are fed forward to a range of extrastriate areas. Thalamocortical and corticocortical connectivity in the visual system can be categorised anatomically as feedforward or feedback (Bullier, 2004). On the basis of this distinction, the different visual areas have been arranged in a structural hierarchy with partly parallel streams

{Felleman:1991km, VanEssen:1983ea, VanEssen:2004vf}. This structural organisation of the visual system shows consistent similarities with functional hierarchies of information processing (Bond, 2004): neurons in early cortical areas are particularly sensitive to basic stimulus properties such as local edge information in the incoming light array and provide the inputs to later areas, which are tuned to increasingly complex stimulus properties. While the general notion of a hierarchical organisation of the visual system has been useful and influential, some degree of caution is warranted given the amount of disagreement about the exact hierarchical level of certain cortical areas (Bullier, 2004; Grill-Spector & Malach, 2004; Hilgetag, O'Neill, & Young, 1996). At present, it seems fair to say that beyond cortical areas V1, V2, and V3 an assignment of a given area to a specific structural or functional level is associated with some uncertainty.

An early and influential attempt to characterise vision computationally can be found in David Marr's seminal work (Marr, 1982). He proposed several levels of computation based on a cascade of filtering stages, starting with the retinal image and progressively building up 3D representations of a visual scene. The computational levels of this model, and of those following in its footsteps, are again hierarchically organised, mirroring the organisation of the bottom-up pathway of the primate visual system. A common confusion in current discussions on top-down processing arises from the assumption that prior information or 'knowledge' that is independent from retinal input can influence visual processing only via top-down modulation. It is typical, that despite being conceptualised as purely bottom-up, computations in models such as Marr's heavily rely on information other than that contained in the retinal image. Inputs that emanate from the environment are structured and exhibit certain regularities (Brunswik & Kamiya, 1953; Geisler, 2007; Simoncelli & Olshausen, 2001). Many organisms exploit this fact by finely tuning their sensory and perceptual systems to the structure of the environmental properties that are relevant to their survival and reproduction. To name just two examples, linear and non-linear receptive field properties of neurons in the early visual system of primates are structured in such a way as to most efficiently code natural environments (Bell & Sejnowski, 1997; Olshausen & Field, 1996; Rao & Ballard, 1999; Schwartz & Simoncelli, 2001) and contour integration processes in humans mirror the contour structures of objects in natural scenes (Geisler & Perry, 2009). These examples illustrate that the structure of biological vision systems reflects relevant properties of an organism's surround. In other words, prior information about the statistical regularities in an organism's environment is embodied in the way in which the visual system processes incoming signals. This kind of restriction on the way in which information can be processed is often referred to as natural constraints. It is widely acknowledged that human vision would be impossible without incorporating prior information in the form of natural constraints because the retinal image drastically underspecifies the relevant aspects of our three-dimensional environment – a phenomenon often referred to as the 'poverty of the stimulus'.

Interestingly, while David Marr is often portrayed as the redeemer of pure bottom-up models, he did not deny the existence of top-down processes (Marr,

1976; 1982). Rather, in his model-building, he used the restriction to bottom-up processes as a heuristic tool to explore how far biological and machine vision could get without top-down processing and, indeed, he thought that such processes are of minor importance in early vision. Today, many vision researchers take the opposite perspective and acknowledge that full recovery of a visual scene under natural conditions is most likely impossible without top-down processing due to difficulties in correct segmentation of objects in the face of occlusion, shadows, motion, reflections, luminance edges and gradients due to lighting, etc. (e.g., (Bullier, 2001; Cavanagh, 2011; Gilbert & Sigman, 2007; Lamme & Roelfsema, 2000; Lee & Mumford, 2003; Roelfsema, 2006; Sillito, Cudeiro, & Jones, 2006). One of the main limitations of pure bottom-up processing is that it cannot combine information from the detailed local analysis of areas such as V1 and V2 with the more global integration of information carried out in higher-level areas such as MT or V4. Neurons in V1 and V2 have small receptive fields; furthermore, lateral connectivity within V1 and V2 is spatially limited and, due to its slow temporal properties, even a cascade of lateral connections could not mediate integration of inputs over larger areas (Angelucci & Bressloff, 2006). These cortical areas therefore provide a very precise representation of local visual properties but are unable to integrate information across long distances in the visual field. Neurons in higher-level areas have larger receptive fields and can do exactly that but lose the fine-grained precision of V1 and V2 neurons. While a consensus on how exactly top-down modulation via feedback connections is linked to bottom-up processing has not yet been reached, it is generally thought to achieve an integration of local and global levels of analysis (Bullier, 2001; 2004; Lamme & Roelfsema, 2000; Sillito et al., 2006).

Top-down processing even at the earliest levels of the visual system has been described in some detail in primates using simultaneous recordings of single- and multi-units at different sites (for reviews see (Bullier, 2001; 2004; Hupe et al., 1998; Sillito et al., 2006). From this work a detailed picture emerges suggesting an important role of top-down influences from V1 onto LGN neurons as well as from MT onto V1 for the perception and segregation of moving objects. In particular, local V1 circuitry is modulated by MT in a stimulus-specific manner so that processing of features represented at the more global analysis are optimised; V1 activity in turn has similar effects on LGN responses. Hence, V1 and MT shape the response properties of cells providing their own inputs. Modulation of LGN and V1 responses by MT are made possible by the fast transmission of information through the magnocellular pathway to MT, allowing feedback to influence how signals relayed through more slowly conducting parvo- and koniocellular-dominated pathways are processed at lower levels (Bullier, 2001; Sillito et al., 2006). Studies on human observers support the notion that top-down influences from the human motion complex MT+/V5 onto V1 are critical for the perception of objects in motion (Pascual-Leone & Walsh, 2001; Silvanto, Cowey, Lavie, & Walsh, 2005; Sterzer, Haynes, & Rees, 2006).

Many further examples of top-down processing via feedback connections come from studies on various aspects of perceptual organization such as figure-ground segregation (e.g. (Lamme, 1995; Self, van Kerkoerle, Supèr, & Roelfsema, 2013),

contour integration (e.g. {Shpaner:2013hy, Altmann:2003wh}, and shape perception (e.g. {Drewes:2016ky, Murray:2004dw}. These studies have been conducted in both animals and humans, and have employed different psychophysical and neurobiological techniques. Additional support comes from work on illusory contour perception (Murray & Herrmann, 2013; Nieder, 2002; Seghier & Vuilleumier, 2006). Certain stimulus configurations induce the perception of a contour that has no physical match in the environment, a so-called illusory contour. Electrophysiological studies that recorded simultaneously at different sites in the primate brain provide evidence suggesting that illusory contours emerge through top-down influences from V2 (and possibly higher cortical areas) onto V1 (Lee & Nguyen, 2001). The functional interpretation of these findings again highlight the importance of integrating information from the fine-grained measurements in V1 with a more global analysis conducted in V2 and higher cortical areas. Studies on human observers using a range of different neurobiological techniques largely support this finding (Murray & Herrmann, 2013; Seghier & Vuilleumier, 2006). The most direct demonstration in humans comes from a study by Wokke and colleageus (2012) who used transcranial magnetic stimulation (TMS) to disrupt perceptual processing in early (V1/V2) and late (LOC) parts of the visual system at different time delays from stimulus onset. Results indicate that all brain areas are involved in processing illusory contours but the critical timing for disruption in the different areas suggests that feedback from LOC onto V1 and V2 plays a critical role.

This is only the tip of the iceberg of the large body of evidence for top-down processing within the visual system. The examples that we have summarised in this section are only meant as an illustration of this body of literature and to show that top-down influences via feedback connections within the visual system are an important part of visual information processing. The function of top-down processing is to allow an integration of local and global levels of analysis, which leads to an optimisation of response properties of the overall system to deal with currently incoming inputs.

IV. Top-down influences from outside the visual system

In recent years, several studies provided evidence to suggest that processing levels that most would consider to be outside the visual system can exert top-down influences on early vision. Here, we will discuss examples of how expectation and memory representations influence processing at early visual stages and we will argue that the function of such top-down effects can be readily understood in terms of the well-established processes discussed in the previous section.

Philosophers, psychologists and neuroscientists often mean very different things when they talk about expectations. At first glance, the concept of expectations seems to be closely related to that of attention (in the sense that changes in expectations often lead to changes in attention and vice versa). Closer analysis, however, reveals that both concepts can and should be kept apart. We will follow

the neuroscientific parlance in what follows and assume that expectation reflects prior information about the probabilistic nature of our environment, whereas attention reflects current motivational relevance (for a review see Summerfield & de Lange, 2014; Summerfield & Egner, 2009). Of course, expectation and attention often coincide in natural viewing situations. Nevertheless, carefully controlled experiments can dissociate their respective contribution to visual processing by independently manipulating probability of stimulus occurrence and task-relevance. For instance, using fMRI and multivariate pattern analysis, it has been demonstrated that the expectation of an upcoming stimulus decreases the overall activity in early visual cortex (V1) but increases its information content (Kok, Jehee, & de Lange, 2012). This pattern of results suggests that expectation of an upcoming stimulus leads to a restructuring of earlier processing mechanisms that sharpens and optimises their response properties in line with the expected stimulus. Importantly, this effect can be dissociated from the task-relevance of the expected features, thereby supporting the notion that the finding is separate from attentional modulation.[3] A range of other studies that employ computational modelling, psychophysical techniques, or neuroimaging come to a similar conclusion (e.g. (Kok, Failing, & de Lange, 2014; Summerfield & Koechlin, 2008; Wyart, Nobre, & Summerfield, 2012). An important question is of course whether predictions of forthcoming sensory events are coded within the visual system. While this is a possibility, current evidence suggests that the brain's memory and executive systems are more likely candidates {Summerfield:2009gb}.

A second set of studies that is of relevance in the current context has focussed on the role of memory representations in shaping early visual processing. Moore and Cavanagh (1998) addressed this question in a psychophysical study, in which they systematically explored two-tone image perception. On first viewing, two-tone images appear like meaningless black and white patches. However, once the observer has viewed the original template image from which the two-tone was created and has therefore received prior information about image content, the visual system is able to bind patches together into a coherent percept. Given that sensory stimulation remains identical before and after the perceptual change, two-tone image perception potentially provides an ideal index of the role of top-down influences from memory representations on perception. Based on detailed experiments that scrutinised the role of various factors in the perception of these stimuli, Moore and Cavanagh (1998) concluded that two-tone image disambiguation requires a top-down processing approach, in which the observer's information processing system follows a global-to-local direction of analysis: Based on memory of semantic content, individual volumetric parts and other details of the objects are recovered.

---

[3] Of course, attention can influence visual processing at all levels. The purpose of this discussion is not to demonstrate that expectations can, but attention cannot influence early visual processing, but rather to separate out influences from expectations and those from attention. The extent to which attention can be considered a top-down influence on perception is a whole different discussion that we will not explore in this paper.

Neuroimaging studies confirm the conclusions from this behavioural work and provide detailed information about the localisation of the measured effects along the visual processing stream. Of particular interest is a study that showed how the perceptual change from meaningless patches to a coherent percept that can be generated with two-tone images is accompanied by changes in the pattern of activity even in early visual areas (V1, V2, V3) (Hsieh, Vul, & Kanwisher, 2010). In particular, the similarity in the activation patterns found in response to the two-tone image and the full template image increases when the two-tone image was perceived as a coherent percept in comparison to when it was seen as meaningless patches. This result confirms the notion that early information processing is reshaped under the instruction of the high-level interpretation of a stimulus. The source of this top-down process are memory-related brain areas such as precuneus and dorsolateral prefrontal cortex (e.g., Dolan et al., 1997; Hegdé & Kersten, 2010).

It is noteworthy that these top-down influences are not specific to artificial stimuli such as two-tone images but can also be measured using pictures of natural scenes, suggesting that top-down influences constitute the typical processing mode under natural viewing conditions. In a recent psychophysical study, Neri (2014) embedded noisy edge elements into natural scenes either in line with the contours of an object or orthogonal to this contour. Using sophisticated reverse-correlation analyses, the perceptual filter of low-level feature detectors was recovered, that observers used to judge whether the embedded edge element was in line with the contour of the object or not. Importantly, filters were measured in two different conditions: The same images were either shown upright or upside down. The rationale of this manipulation is that early visual areas do not care about this manipulation; these images provide the same input to neurons in the early parts of the visual system. However, the extraction of semantic content is disrupted in images that are presented upside down. The study found that the tuning of perceptual filters was influenced by the extent to which semantic information could be extracted from the images. This suggests that the global, semantic content of a visual scene can reshape the properties of local information processing units early in the visual system, a conclusion that is supported by a computational model of these effects. It is critical to note that that this study explicitly controlled for effects of attention and found the facilitatory effect of top-down modulation by semantic content to be orthogonal to that of attentional modulation.

The influences of expectation and memory on perception that we have mentioned above are presumably closely related. In fact, some proposals view both processes as part of the same top-down machinery (Bar, 2009). Regardless of whether this is the case, it is noteworthy that the expectation- and memory-induced changes in the responsiveness of signal-selective units that most likely underlies sharpening of representations in early visual areas is reminiscent of the effects of top-down processing *within* the visual system that we discussed in the previous section. In other words, similarly to how influences from high- onto low-level *perceptual* processing provide a means of combining global and local levels of perceptual analyses, top-down influences of expectation and memory onto perception allow an observer to integrate relevant information at an even

more global level with sensory inputs. The ultimate aim of top-down influences in both instances is the optimisation of information processing at the system level. This strategy does not come without cost. Top-down influences from expectation and memory representations have been demonstrated to be of clinical significance (Schmack et al., 2013; Teufel et al., 2015), where an undue reliance on this type of processing can increase the risk of a loss of contact with reality characteristic of psychosis.

V. Conclusions
Humans and nonhuman primates live in a highly complex and ever-changing visual world. In order to deal with the resultant uncertainty and to generate adaptive representations that can guide successful behaviour, visual information processing is highly interactive with an important role for top-down influences from higher-level representations onto lower-level perceptual processing. While some of these higher-level representations can safely be considered to be perceptual – and the resultant top-down influence can be considered to be located *within* the visual system – we argue, that current evidence suggests the existence of top-down influences that some would consider to cross the perception/cognition divide. Independent of the source, the purpose of all of these instances of top-down processing seems to be an adjustment of low-level, local circuitry to current visual inputs. This optimisation leads to sparse representations that enable both fine-scale resolution of sensory input and large-scale integration through higher-level representations.[4]

---

[4] It is an important question how top-down influences on perceptual processing (of the kind we focused on) relate to philosophical questions about how 'cognitive penetrability' poses a problem for some accounts of perceptual justification (see Siegel 2011). While we cannot address the complex issue here for lack of space, see Nanay forthcoming for a possible account of the epistemic consequences of top-down influences of the kind analysed in this paper.

References:

Angelucci, A., & Bressloff, P. C. (2006). Contribution of feedforward, lateral and feedback connections to the classical receptive field center and extra-classical receptive field surround of primate V1 neurons. *Progress in Brain Research*, *154*, 93–120. http://doi.org/10.1016/S0079-6123(06)54005-1

Bar, M. (2009). The proactive brain: memory for predictions. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*(1521), 1235–1243. http://doi.org/10.1098/rstb.2008.0310

Bayne, T. (2009). Perception and the reach of phenomenal content. *Philosophical Quarterly* 59: 385-404.

Bell, A. J., & Sejnowski, T. J. (1997). The "independent components" of natural scenes are edge filters. *Vision Research*, *37*(23), 3327–3338. http://doi.org/10.1016/S0042-6989(97)00121-1

Bond, A. H. (2004). An information-processing analysis of the functional architecture of the primate neocortex. *Journal of Theoretical Biology*, *227*(1), 51–79. http://doi.org/10.1016/j.jtbi.2003.10.005

Brindley, G. S. (1960). Physiology of the Retina and Visual Pathway. London: Edward Arnold.

Brunswik, E., & Kamiya, J. (1953). Ecological Cue-Validity of "Proximity" and of Other Gestalt Factors. *The American Journal of Psychology*, *66*(1), 20. http://doi.org/10.2307/1417965

Bullier, J. (2001). Integrated model of visual processing. *Brain Research Reviews*, *36*(2-3), 96–107. http://doi.org/10.1016/S0165-0173(01)00085-6

Bullier, J. (2004). Communications between Cortical Areas of the Visual System. In L. M. Chalupa & J. S. Werner (Eds.), *The Visual Neurosciences* (pp. 522–540). Cambridge, Masachusetts: The MIT Press.

Cavanagh, P. (2011). Visual cognition. *Vision Research*, *51*(13), 1538–1551. http://doi.org/10.1016/j.visres.2011.01.015

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, *36*(3), 181–253. http://doi.org/10.1017/S0140525X12000477

Dolan, R. J., Fink, G. R., Rolls, E., Booth, M., Holmes, A., Frackowiak, R. S. J., & Friston, K. J. (1997). How the brain learns to see objects and faces in an impoverished context. *Nature*, *389*(6651), 596–599. http://doi.org/10.1038/39309

Dunham, Y. et al. (2008). The development of implicit intergroup cognition, *Trends in Cognitive Sciences*, 12: 248-253.

Felleman, D. J., & Van Essen, D. C. (1991). Distributed Hierarchical Processing in the Primate Cerebral Cortex. *Cerebral Cortex*, *1*(1), 1–47. http://doi.org/10.1093/cercor/1.1.1

Firestone, C., & Scholl, B. J. (2016). There are no top-down influences on perception. *Behavioral and Brain Sciences*, forthcoming

Fletcher, P. C., & Frith, C. D. (2009). Perceiving is believing: a Bayesian approach to explaining the positive symptoms of schizophrenia. *Nature Reviews Neuroscience*, *10*(1), 48–58. http://doi.org/10.1038/nrn2536

Geisler, W. S. (2007). Visual Perception and the Statistical Properties of Natural Scenes. *Dx.Doi.org*, *59*(1), 167–192. http://doi.org/10.1146/annurev.psych.58.110405.085632

Geisler, W. S., & Perry, J. S. (2009). Contour statistics in natural images: Grouping across occlusions. *Visual Neuroscience*, *26*, 109–121.

Gilbert, C. D., & Sigman, M. (2007). Brain States: Top-Down Influences in Sensory Processing. *Neuron*, *54*(5), 677–696. http://doi.org/10.1016/j.neuron.2007.05.019

Greenwald, A. G., & Banaji, M.R. (1995). Implicit social cognition. *Psychological Review*, 102: 4-27.

Grill-Spector, K., & Malach, R. (2004). The human visual cortex. *Annual Review of Neuroscience*, *27*(1), 649–677. http://doi.org/10.1146/annurev.neuro.27.070203.144220

Haggard, P., Clark, S., & Kalogeras, J. (2002). Voluntary action and conscious awareness. *Nature Neuroscience, 5(4)*, 382–385.

Hegdé, J., & Kersten, D. (2010). A Link between Visual Disambiguation and Visual Memory. *Journal of Neuroscience*, *30*(45), 15124–15133. http://doi.org/10.1523/JNEUROSCI.4415-09.2010

Hilgetag, C. C., O'Neill, M. A., & Young, M. P. (1996). Indeterminate organization of the visual system. *Science*, *271*, 776–777.

Hsieh, P. J., Vul, E., & Kanwisher, N. (2010). Recognition Alters the Spatial Pattern of fMRI Activation in Early Retinotopic Cortex. *Journal of Neurophysiology*, *103*(3), 1501–1507. http://doi.org/10.1152/jn.00812.2009

Hupe, 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. http://doi.org/10.1038/29537

Jogan, M., & Stocker, A. A. (2014). A new two-alternative forced choice method for the unbiased characterization of perceptual bias and discriminability. *Journal of Vision*, *14*(3), 20–20. http://doi.org/10.1167/14.3.20

Kahneman, D. & Tversky, A. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology,* 5, 207-233.

Katzner, S., & Weigelt, S. (2013). Visual cortical networks: of mice and men. *Current Opinion in Neurobiology*, *23*(2), 202–206. http://doi.org/10.1016/j.conb.2013.01.019

Kingdom, F., & Prins, N. (2009). Psychophysics: A Practical Introduction. Amsterdam: Elsevier.

Kok, P., Failing, M. F., & de Lange, F. P. (2014). Prior Expectations Evoke Stimulus Templates in the Primary Visual Cortex. *Journal of Cognitive Neuroscience*, *26*(7), 1546–1554. http://doi.org/10.1162/jocn_a_00562

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. http://doi.org/10.1016/j.neuron.2012.04.034

Kriegel, U. (2007). The phenomenologically manifest. *Phenomenology and the Cognitive Sciences 6*, 115-136.

Lamme, V. A. (1995). The neurophysiology of figure-ground segregation in primary visual cortex. *Journal of Neuroscience*, *15*(2), 1605–1615.

Lamme, V. A. F., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences*, *23*(11), 571–579. http://doi.org/10.1016/S0166-2236(00)01657-X

Lee, T. S., & Mumford, D. (2003). Hierarchical Bayesian inference in the visual cortex. *Journal of the Optical Society of America A - Optics Image Science and*

Vision, *20*(7), 1434–1448. http://doi.org/10.1364/JOSAA.20.001434

Lee, T. S., & Nguyen, M. (2001). Dynamics of subjective contour formation in the early visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *98*(4), 1907–1911. http://doi.org/10.1073/pnas.031579998

Lupyan, G., Thompson-Schill, S. L., & Swingley, D. (2010). Conceptual penetration of visual processing. *Psychological Science*, *21*(5), 682–691.

Lupyan, G. (2015). "Cognitive Penetrability of Perception in the Age of Prediction: Predictive Systems are Penetrable Systems." *Review of Philosophy and Psychology*: 1-23.

Macpherson, F. 2012 Cognitive penetration of colour experience. *Philosophy and Phenomenological Research* 84: 24-62

Marr, D. (1976). Early Processing of Visual Information. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *275*(942), 483–524. http://doi.org/10.1098/rstb.1976.0090

Marr, D. (1982). Vision. Cambridge, Massachusetts: The MIT Press.

Masrour, F. (2011). Is perceptual phenomenology thin? *Philosophy and Phenomenological Research* 83: 366-397.

Moore, C., & Cavanagh, P. (1998). Recovery of 3D volume from 2-tone images of novel objects. *Cognition*, *67*(1-2), 45–71. http://doi.org/10.1016/S0010-0277(98)00014-6

Morgan, M. J., Melmoth, D., & Solomon, J. A. (2013). Linking hypotheses underlying Class A and Class B methods. *Visual Neuroscience*, *30*(5-6), 197–206. http://doi.org/10.1017/S095252381300045X

Morgan, M., Dillenburger, B., Raphael, S., & Solomon, J. A. (2012). Observers can voluntarily shift their psychometric functions without losing sensitivity. *Attention, Perception, & Psychophysics*, *74*(1), 185–193. http://doi.org/10.3758/s13414-011-0222-7

Murray, M. M., & Herrmann, C. S. (2013). Illusory contours: a window onto the neurophysiology of constructing perception. *Trends in Cognitive Sciences*, *17*(9), 471–481. http://doi.org/10.1016/j.tics.2013.07.004

Neri, P. (2014). Semantic Control of Feature Extraction from Natural Scenes. *Journal of Neuroscience*, *34*(6), 2374–2388. http://doi.org/10.1523/JNEUROSCI.1755-13.2014

Nanay, B. (2011a). Do we perceive apples as edible? *Pacific Philosophical Quarterly* 92: 305-322.

Nanay, B. (2012). Perceptual Phenomenology. *Philosophical Perspectives* 26: 235-246

Nanay, B. (2013). *Between Perception and Action*. Oxford: Oxford University Press.

Nieder, A. (2002). Seeing more than meets the eye: processing of illusory contours in animals. *Journal of Comparative Physiology*, *188*(4), 249–260. http://doi.org/10.1007/s00359-002-0306-x

Nisbett, R. E. & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review 84*: 231-259.

Olshausen, B. A., & Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, *381*, 607–609.

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. http://doi.org/10.1126/science.1057099

Pylyshyn, Z. (1999). Is vision continuous with cognition?: The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, *22*(03), 341–365.

Rao, R., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, *2*(1), 79–87.

Roelfsema, P. R. (2006). Cortical algorithms for perceptual grouping. *Annual Review of Neuroscience*, *29*(1), 203–227. http://doi.org/10.1146/annurev.neuro.29.051605.112939

Schmack, K., de Castro, A. G.-C., Rothkirch, M., Sekutowicz, M., Rössler, H., Haynes, J.-D., et al. (2013). Delusions and the Role of Beliefs in Perceptual Inference. *Journal of Neuroscience*, *33*(34), 13701–13712. http://doi.org/10.1523/JNEUROSCI.1778-13.2013

Schwartz, O., & Simoncelli, E. P. (2001). Natural signal statistics and sensory gain control. *Nature Neuroscience*, *4*(8), 819–825. http://doi.org/10.1038/90526

Schwitzgebel, E. (2008). The Unreliability of Naïve Introspection. *Philosophical Review,* 117, 245-273.

Seghier, M. L., & Vuilleumier, P. (2006). Functional neuroimaging findings on the human perception of illusory contours. *Neuroscience and Biobehavioral Reviews*, *30*(5), 595–612. http://doi.org/10.1016/j.neubiorev.2005.11.002

Self, M. W., van Kerkoerle, T., Supèr, H., & Roelfsema, P. R. (2013). Distinct Roles of the Cortical Layers of Area V1 in Figure-Ground Segregation. *Current Biology*, *23*(21), 2121–2129.

Siegel, S. (2007). How can we discover the contents of experience? *Southern Journal of Philosophy* (Supp), *45*, 127-142.

Siegel, S. (2011). Cognitive penetrability and perceptual justification. *Nous*, *46*, 201-222.

Siewert, C. (2002). Is visual experience rich or poor? *Journal of Consciousness Studies*, *9*, 131-140.

Sillito, A. M., Cudeiro, J., & Jones, H. E. (2006). Always returning: feedback and sensory processing in visual cortex and thalamus. *Trends in Neurosciences*, *29*(6), 307–316. http://doi.org/10.1016/j.tins.2006.05.001

Silvanto, J., Cowey, A., Lavie, N., & Walsh, V. (2005). Striate cortex (V1) activity gates awareness of motion. *Nature Neuroscience*, *8*(2), 143–144. http://doi.org/10.1038/nn1379

Simoncelli, E. P., & Olshausen, B. A. (2001). Natural Image Statistics and Neural Representations. *Annual Review of Neuroscience*, *24*(1), 1193–1216. http://doi.org/10.1146/annurev.neuro.24.1.1193

Spener, M. and Bayne, T. (2010). Introspective Humility. *Philosophical Issues*, *20*, 1-22.

Sterzer, P., Haynes, J.-D., & Rees, G. (2006). Primary visual cortex activation on the path of apparent motion is mediated by feedback from hMT+/V5. *NeuroImage*, *32*(3), 1308–1316.

Stokes, D. (2012). Perceiving and Desiring: A new look at the cognitive penetrability of experience. *Philosophical Studies, 158*, 479-92.

Storrs, K. R. (2015). Are high-level aftereffects perceptual? *Frontiers in Psychology*, *6*, 373. http://doi.org/10.3389/fpsyg.2015.00157

Summerfield, C., & de Lange, F. P. (2014). Expectation in perceptual decision

making: neural and computational mechanisms. *Nature Reviews Neuroscience*, *15*(11), 745–756. http://doi.org/10.1038/nrn3838

Summerfield, C., & Egner, T. (2009). Expectation (and attention) in visual cognition. *Trends in Cognitive Sciences*, *13*(9), 403–409. http://doi.org/10.1016/j.tics.2009.06.003

Summerfield, C., & Koechlin, E. (2008). A Neural Representation of Prior Information during Perceptual Inference. *Neuron*, *59*(2), 336–347. http://doi.org/10.1016/j.neuron.2008.05.021

Teufel, C., Subramaniam, N., Dobler, V., Perez, J., Finnemann, J., Mehta, P. R., et al. (2015). Shift toward prior knowledge confers a perceptual advantage in early psychosis and psychosis-prone healthy individuals. *Proceedings of the National Academy of Sciences of the United States of America*, *112*(43), 13401–13406. http://doi.org/10.1073/pnas.1503916112

Tversky, A. and Daniel, K. (1981). The Framing of Decisions and the Psychology of Choice. *Science*, *211*, 453-458.

Van Essen, D. C. (2004). Organization of Visual Areas in Macaque and Human Cerebral Cortex. In L. M. Chalupa & J. S. Werner (Eds.), *The Visual Neurosciences* (pp. 507–521). Cambridge, Masachusetts: The MIT Press.

Van Essen, D. C., & Maunsell, J. H. R. (1983). Hierarchical organization and functional streams in the visual cortex. *Trends in Neurosciences*, *6*(9), 370–375. http://doi.org/10.1016/0166-2236(83)90167-4

Vetter, P. and A. Newen (2014). Varieties of cognitive penetration in visual perception. *Consciousness & Cognition* 27: 62-75.

Wegner, D. M. (2002). *The illusion of conscious will*. Cambridge: MIT Press.

Wegner, D. M., Sparrow, B., & Winerman, L. (2004). Vicarious agency: Experiencing control over the movements of others. *Journal of Personality and Social Psychology 86*: 838-848.

Williams, L. E. and Bargh, J. A. (2008). Experiencing Physical Warmth Promotes Interpersonal Warmth. *Science*, *322*, 606-607.

Wokke, M. E., Vandenbroucke, A. R. E., Scholte, H. S., & Lamme, V. A. F. (2012). Confuse Your Illusion: Feedback to Early Visual Cortex Contributes to Perceptual Completion. *Psychological Science*, *24*(1), 0956797612449175–71. http://doi.org/10.1177/0956797612449175

Wyart, V., Nobre, A. C., & Summerfield, C. (2012). Dissociable prior influences of signal probability and relevance on visual contrast sensitivity. *Proceedings of the National Academy of Sciences of the United States of America*, *109*(9), 3593–3598.

Zhong, C. and Liljenquis, K. (2006). Washing Away Your Sins: Threatened Morality and Physical Cleansing. *Science*, *313*, 1451-1452.