

Perceptual Integration, Modularity, and Cognitive Penetration*

Daniel Burnston and Jonathan Cohen[†]

Imagine a computational system with a proprietary ... database. Imagine that this device operates to map its characteristic inputs onto its characteristic outputs ... and that, in the course of doing so, its informational resources are restricted to what its proprietary database contains. That is, the system is "encapsulated" with respect to information that is not in its database. ... That's what I mean by a module. In my view, it's informational encapsulation, however achieved, that's at the heart of modularity (Fodor, 2001, 63).

On the best evidence, there seems to be significant interaction between information sources, at all levels of perception: for better or worse, the picture of perception as a collection of independent, non-interacting, exclusively bottom-up feature extractors must be replaced by a picture of perception that more centrally incorporates informational interaction. The question we want to raise in this paper is what impact this lesson should have for our thinking about claims about the modularity and the cognitive penetrability of perception.

Many authors have taken the news that perception involves substantial informational integration to require seriously weakening or abandoning ideas about modularity and impenetrability. We believe this is an overreaction. For, we claim, an integrative conception of perception is compatible with robust forms of both modularity and impenetrability. That said, we believe that standard conceptions of the criteria for individuating modules sit uneasily with the integrative view of perception — that the standard view involves associating processes with univocal informational functions, and then looking for exchange of information (or the lack of it) between distinct processes, in a way that is hard to reconcile with the integrative view. As such, adopting the latter will require some rethinking of our standards for drawing modular boundaries. In what follows we'll be offering just such a reconception of the notion of modularity in terms of what Fodor (1983) calls 'isotropy'. We claim that the isotropy-based criterion both comports better with the integrative view and better reflects the underlying purposes for which the notion of modularity was introduced, so is well worth adopting. Moreover, the isotropy-based condition suggests ways in which the usual picture, on which the modular/non-modular

*This paper is entirely collaborative; authors are listed in alphabetical order.

[†]Department of Philosophy, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0119, dburnsto@ucsd.edu, joncohen@aardvark.ucsd.edu

distinction coincides exactly with the perception/cognition distinction, is oversimple.

Here is our plan. We'll begin by reviewing, briefly, the case for an integrative conception of perception (§1). Next, we'll turn to the literature on modularity, and argue that both foes and ostensible friends of modularity have overestimated the threat posed to the notion by facts about informational integration (§2). We'll then pose a dilemma about the formulation of modularity in an integrativist framework, and suggest abandoning some of the standard criteria for modularity in favor of an alternative based on the notion of isotropy (§3). Then we'll show how the resulting, integrativist-friendly reformulation of modularity applies in particular cases, how it can be used to make fruitful distinctions with respect to both perception and cognition, and what this tells us about cognitive penetrability and perception (§4). Finally, we'll conclude (§5).

We hope that our discussion will bring out more clearly what is at stake in debates about perception, modularity and cognitive penetration, and will offer a new and improved way of thinking about these topics in light of the lessons of perceptual integration.

1 The Unencapsulation of Perception

1.1 A (Too) Simple View

Consider the following logically possible, coherent, cartoon conception of perception, which we can call *the non-integrative view*.¹ On this view, perception consists of a series of separate, dedicated, and informationally encapsulated modules, each devoted to the extraction of a single feature of the distal layout from the impinging energy array. Perhaps there is, for example, a shape extractor that produces as output a representation with the content that the object of perception is square; a color extractor that produces as output a representation with the content that the object of perception is green; a motion extractor that produces as output a representation with the content that the object of perception is moving along a particular path with a particular velocity; and so on. Once these extractors have done their jobs, they hand off their output to a higher level module that parses the scene into visual objects on the basis of information from the feature extractors about how features are distributed throughout the scene.² The non-integrativist's stipulation

¹We emphasize that the view we are sketching here is a caricature that should not be identified with the views of any actual thinker (though of course there may be points of overlap). We offer this deliberately overdrawn cartoon in order to locate its vulnerabilities easily (§1.2) and motivate a more sophisticated alternative (§1.3).

²Here and below we use the term 'object' in a thin sense that picks out (roughly) bounded and connected things that trace out continuous paths in spacetime, but that needn't retain their kind affiliations to count as persisting over time. Minimally, we want it to be true of objects that they can be distal, that they can bear features that perceptual systems can be sensitive to, that they can serve as the loci for binding of feature representations, and that they can be served up by perception qua

that the feature extractors are mutually informationally encapsulated is the stipulation that each does its computational work independently — that the inputs to the computations carried out by the shape extractor do not include the outputs of the color extractor (as it might be), nor vice versa. The view does allow for informational/computational dependence of some higher level stages of perceptual representation — say, for object representation — on the output of the feature extractors. But, crucially, the non-integrativist views this dependence as asymmetric: the story would be that object representations are computed from representations about the exemplification of features, but not vice versa.

Besides being an attractively simple story about the general architecture of perception, the non-integrative view has the virtue of smoothing the way to an extremely simple and direct verdict on the question of whether perception is cognitively penetrated. For, because the view is constructed so as to guarantee that individual perceptual faculties are encapsulated from *all* other representations, it follows immediately that they are encapsulated from cognitive representations in particular. Moreover, as have recounted elsewhere (Burnston and Cohen, 2013), the non-integrative view has played an important role in thinking about perception in philosophy, computational psychology, and neuroscience.

Given all these merits of the non-integrative view, it is worth asking: is it true?

1.2 Perceptual Interaction, Within and Between Modalities

Alas, it turns out, it is not. There is abundant and converging evidence, drawn from many levels of organization in many different perceptual systems, that the architecture of perceptual computation is significantly more interactive than the non-integrative view allows, both within and between sensory modalities.³

A first kind of evidence of intramodal interactions comes from Field *et al.* (2009), who show that rod photoreceptors, which ordinarily function under scotopic conditions, provide input to blue/yellow opponent pathways in the macaque retina. A related idea, suggested (speculatively) by Broackes (2009), is that human deuteranopes employ dynamic input from their intact S-cone receptors to extract color information that is ordinarily subserved in trichromats by the output of L-cones or M-cones. If correct, these proposals

the referents of demonstrative thought and language on at least some occasions. Beyond that, we invite readers to rely on their own preferred understanding of what thin objects amount to.

³A further reason for dissatisfaction with the non-integrative view revolves around its insistence that the psychological representation of objects is always posterior to the psychological representation of features. As we have argued elsewhere (Burnston and Cohen, 2013), there are reasons to think that perception can represent the presence and location of objects in ways that outstrip the evidence drawn from its representations of exemplified features; and there are cases in which it seems that the perceptual representation of features depends on the perceptual representations of objects, contrary to the non-integrativist's preferred direction of representational dependency. We put these complications to one side in the following.

reveal a level of unencapsulation at the photoreceptor level: they suggest that blue/yellow extractors are not informationally encapsulated from rod scotopic output, and that red/green extractors are not encapsulated from S-cone output.⁴

Nor do interactions within color vision stop with retinal processing: a variety of interactive effects occur in the process of assigning colors to surfaces, and in the computation of color constancies (see Cohen (2012) for discussion). Once again, if there is this sort of integration, then this tells against the view that there is any single informational pathway for the extraction of surface color.

Further examples come from a variety of perceptual illusions, which, when manipulated experimentally, can tease apart inputs that affect perceptual processing of scene elements. A particularly revealing set of examples comes from the study of apparent motion.⁵ A variety of results show that apparent motion is affected by both luminance and contrast, and moreover that the speed of motion, its direction, and motion-related object perception are all affected, depending on the perceptual situation. For instance, Anstis (2003b, 2004) demonstrates that the speed of a square patch moving at a constant velocity against a black/white grated background seems to change as the leading edge contrast changes. When the contrast is low, the square appears to slow down; when contrast is increased, it appears to speed up. When a square or diamond is moved diagonally across a grated or plaid surface, the different speed perceptions produced by the different contrasts of its sides (i.e., against the white or the black background) modify its perceived direction, so that even an object moving on a straight trajectory will appear to swerve back and forth. Similarly, in the “crossover illusion,” subjects are shown four rectangular shapes in the shape of a larger rectangle, where the top rectangles are ordered white/black and the bottom are ordered black/white (or vice versa). Each rectangle then repeatedly switches with the one horizontally adjacent to it. In this setup, subjects perceive one (black or white) object moving diagonally; but, crucially, which color the object is seen as exemplifying is determined by which has the higher contrast compared to the background. Further, in these illusions both contrast and overall luminance can affect the strength of the effect (measured by time until breakdown; Anstis (2003a)). Overall, these

⁴One might object that counting these effects (and other cases discussed in this section) as encapsulation failures is too liberal, and that, on the resulting, highly permissive view, every input to a process (even its canonical ones) will count as encapsulation failures for that process.

This worry is well taken: we do not wish to endorse an excessively rigid construal of the canonical input domains for particular processes, relative to which encapsulation will turn out to be trivially impossible. In a way, however, this is just the point we are making with this and our other examples: the non-integrative view of §1.1 itself sets over-rigid boundaries for particular processes, with the result that there turn out to be any number of influences coming from outside them. Thus, our point is that these cases are counterexamples to unencapsulation *given the way of drawing boundaries accepted by the non-integrativist*. It is compatible with this point that we should decide not to accept the non-integrativist’s way of associating domains and inputs at the end of the day. (Thanks to Athanasios Raftopoulos for urging the need for clarification on this point.)

⁵Similar conclusions can be drawn from illusions involving, for example, luminance and color, or from findings concerning illusory contours.

studies suggest that motion perception draws on temporal sequences of both form and luminance information.⁶ Of course, this is just to deny that there is any single, informationally encapsulated pathway for the extraction of motion.

Multi-modal interaction effects are also common at all perceptual levels. Thus, there is neuroscientific evidence of multi-modal effects at both the sub-cortical (Stein and Stanford, 2008) and primary sensory cortical (Kayser and Logothetis, 2007) levels, as well as in areas of temporal cortex implicated in categorical perception (James *et al.*, 2002).⁷ Moreover, ongoing investigation into such well-known effects as the flash-lag illusion (MacKay, 1958; Nijhawan, 1994) and the ventriloquist illusion (Pick *et al.*, 1969; Vroomen and de Gelder, 2000, 2004) have shown not only that information from visual and auditory modalities can interpenetrate, but that such intermodal interactions can have reliable and positive effects on perceptual learning (Shams *et al.*, 2011). Similarly, the cutaneous rabbit illusion (Geldard and Sherrick, 1972) and other cases of sensory substitution (Bach y Rita *et al.*, 1969) are standardly taken to show interaction between tactile and visual pathways that would preclude their mutual encapsulation.⁸

In short, while there remain many unanswered questions about the nature of these interactions, the prevalence and range of interactive effects within and across modalities is hard to square with the non-integrativist's picture of independent and mutually informationally encapsulated pathways for the extraction of individual features.

1.3 An Integrative Alternative

Given the sorts of findings discussed in §1.2, we need a way of thinking of perception as much more extensively integrated and interactive than the non-integrative view would allow.

In other work (Burnston and Cohen, 2013), we have offered such a conception — henceforth, the integrative view — and argued that is supported by a wide range of broadly empirical considerations. In outline, the integrative view represents the perceptual system as recovering structure about the distal world from an initially unstructured transduced signal by deploying a variety of techniques, or “perceptual strategies,” each operating on a characteristic range of information, and each delivering a working model of the arrangement of the particular features and objects that fall within its domain. As the process

⁶This lesson is reinforced by related studies of visual form detection, which have shown that subjects can accurately compare (without attentional switching) several different kinds of form definition, including motion-defined, luminance-defined, and cyclopean (stereoscopic) form, even on opposite sides of irrelevant distracter forms (Kohly and Regan, 2002).

⁷While there is no consensus on the correct interpretation of these effects, they put significant pressure on the traditional view of unimodal cortices that are combined in higher-level “convergence” or “association” areas (Mesulam, 1998; Meyer and Damasio, 2009).

⁸For discussion of such multimodal interactions and their general significance for philosophy of perception, see O'Callaghan (2013).

continues, perception integrates the different working models to produce a net perceptual representation as a kind of package deal.⁹

Among its other benefits, and as its name suggests, the integrative view is especially well-suited to accommodate the sorts of perceptual integration described in §1.2. Thus, for example, the integrative view easily accommodates the apparent fact that the assignment of colors to surfaces (etc.) involves integrating the output of a range of different computational techniques, each of which can be thought of as a perceptual strategy that delivers a working model from a canonical range of inputs, and whose output working model must be integrated with those of whatever other perceptual strategies have outputs that speak to the description of distal color. Similarly, the integrative view makes sense of the possibility of interaction between perceptual strategies directed at distinct features within or between modalities.¹⁰

Of course, on the integrative view, perception must, at some stage, somehow combine the outputs of multiple perceptual strategies, possibly operating over distinct spatiotemporal scales. Presumably the goal of such integrations is to arrive at a value that is consistent with (or, as consistent as possible with) the range of evidence supplied by the different perceptual strategies at work. And, needless to say, different strategies may deliver inconsistent verdicts (particularly, but not exclusively, in laboratory settings designed specifically to pull them apart). When this happens, the perceptual system must have some way of managing the conflicts. And there is evidence of exactly the sorts of conflict management strategies one might expect. Thus, for example, sometimes perception selectively prefers one strategy over others (e.g., as in our reaction to the Ames room, where perspective information about depth dominates over our expectations about the relative size of objects). In other cases perception combines the information from distinct perceptual strategies in some more complex way, by sending one or more of the strategies back to run on the input again in the hope of obtaining new answers that won't conflict, or by giving up (as, for example, we do with impossible Escher figures). The observation of such effects further supports the view that "the" perceptual representation of the world is in fact an interaction effect arising from the integration of multiple perceptual strategies.

⁹Caveat: Though, in what follows, we present the integrative view as an alternative general architecture for perception, we don't want to commit to the claim that it is a correct description of *every* process occurring in perception. While we find the evidence supporting it to be wide-ranging and compelling, it is obviously an empirical question, not to be prejudged from the armchair, just how much of perception can be understood in terms of perceptual strategies and their integration.

¹⁰Indeed, the view also accommodates the idea that the input to a perceptual strategy may be at least partly endogenous, and may encode "internalized regularities" about the environments in which they are employed. On this general explanatory strategy, which is now fairly ubiquitous in theorizing about perception, see Shepard (2001); for representative instances, see, for example, Ullman (1979); Lotto and Purves (2001); Purves *et al.* (1999); Yuille and Kersten (2006).

2 Modularity Redux?

If the non-integrative view were the only interesting theoretical motivation for the position that aspects of perception are impenetrable/modular, then learning that the former view is false would make the position (at least as classically construed) look pretty hopeless. And, as it happens, many thinkers have endorsed versions of this reasoning (though the responses they have proposed to it have differed). We believe these conclusions are too quick, and that there remain interesting senses in which impenetrability/modularity remain very much on the table. However, before we can spell this out (§§3–4), in the current section we want to urge that both foes and friends of modularity have misconstrued the lessons integration holds for the notion.

On one side, several foes of modularity (particularly those who have focused on the details of cognitive and neural processing) have taken the evidence of interactions between cognitive/perceptual processes to show that modularity is a lost cause. Thus, after discussing both neural plasticity and the sharing of information between neural systems, Buller and Hardcastle (2000) conclude that “there is little sense in which we can maintain that even our most basic processes are distinct in the way that [the] modularity hypothesis implies” (313). Similarly, Prinz (2006) cites a wide range of interactive effects in perceptual and cognitive processes (including some of the multi-modal effects we mentioned in §1.2) and concludes that virtually none of cognition — including early perception — is modular.

On the other side, ostensible friends of modularity (in particular, defenders of massive modularity or evolutionary psychology more broadly speaking) have often taken evidence of the kinds of integration discussed in §1 as reason for weakening or rejecting encapsulation as a criterion of modularity.¹¹ Effectively, these theorists accept with the anti-modularists the idea that the evidence of informational integration refutes modularity *qua* classically conceived, and go on to “save” modularity by replacing it with something weaker. Thus, for example, Coltheart (1999) abandons encapsulation in favor of a loose notion of domain specificity. Sperber (2005) hopes to save modularity by distinguishing senses of domain-specificity relevant for understanding function from those that might be affected by interaction. Carruthers (2006) makes a similar move in offering a “weak” sense of encapsulation that can maintain functional specificity despite interaction effects from other processes. And Barrett and Kurzban (2006) distinguish between the information a module has “access” to and the information in processes.¹²

We suggest that both of these reactions are unjustified.

¹¹Alternatively, such theorists have sometimes attempted to save modularity by playing down the importance of integration in (early) perception, and thereby have characterized perception in terms much closer to the feature extractors we caricatured in §1.1 than to the complex and integrative strategies we propose in §1.3 (Barrett and Kurzban, 2006; Samuels, 2006).

¹²Indeed, some have held that giving defining characteristics of modularity at all obscures the very point of functional analysis of mental processing (Barrett and Kurzban, 2006; Sperber, 2005).

Consider the pro-modularists first. One reason for dissatisfaction with the defenses they have proposed for modularity is that the weakened notions of domain specificity and weak encapsulation on which they rest have so far remained somewhat loose and metaphorical; indeed, we are sympathetic with the complaint that they may lapse into triviality (Samuels, 2006). Second, and more importantly in our view, weakened notions of modularity sacrifice what we take to have been one of the important hallmarks of Fodor's original account of the notion — viz., the advantages modularity carries for purposes of implementing a physically/computationally tractable theory of cognition. Fodor (1983) famously argued that we could understand how cognitive processing might possibly work only to the extent that we could implement it in modular terms. Whether or not Fodor's overall pessimism is justified, the conditions he places on modularity strike us as a potentially useful tool in our collective attempts to understand cognition and perception. Correlatively, giving up these conditions, even to the point of abstracting away from what is known about actual processing altogether, threatens to drain much of the interest from the notion of modularity itself.¹³

As against the anti-modularists, we suggest that interactions between different types of information within perception do not inevitably conflict with modularity. Crucially, and as we have pointed out elsewhere (Burnston and Cohen, 2013), the very notion of modularity makes sense only relative to a basic functional and informational vocabulary specifying what counts as inside and outside a module. Only relative to such a specification of boundaries does it make sense to think of a particular sort of information exchange as a violation of modularity. Now, if you specify boundaries in a vocabulary tied to individual features, then featural integration in perception (e.g., of the sorts reviewed in §1.2) will look like boundary-crossing, hence like modularity violation. But that vocabulary for drawing the boundaries is not mandatory. Indeed, given the extensive evidence of featural integration in perception, we suggest that the perceptual strategies implicated in perceptual processing make for better, and more empirically-motivated, candidates for the kinds in terms of which our theories should be cast. But relative to a (in our view, better motivated) vocabulary that locates modular boundaries in terms of perceptual strategies, instances of informational integration in perceptual processing need not count as crossing a modular boundary. For if we can think of each instance of informational integration as occurring within a single perceptual strategy, we can count it as occurring within the boundaries of a single module.¹⁴

¹³A further concern is that several of these weakened modularity views rely on a not obviously principled distinction between the information proprietary to a cognitive/neural process and the information that is secondary to it. Proponents have sometimes attempted to meet this worry by appeal to evolutionary teleology or neurophysiological evidence. Suffice to say that it is possible to doubt whether these sorts of evidence can sustain the burden of compelling controversial views of cognitive architecture (Young *et al.*, 2000; Buller, 2005).

¹⁴Viewed from this perspective, the appeal to integrative effects as evidence against modularity is a strange argumentative strategy; for it accepts the basic classificatory vocabulary of a subset of modular views (e.g., dedicated feature extraction views) only to deny it by citing an overwhelming abundance of evidence against it.

This shows that, on our preferred way of thinking of modularity, the evidence of informational integration is at least possibly consistent with the claim that perceptual systems are modular. That said, it doesn't follow immediately from what we have said that modularity is *true*: this will depend on the (empirical) questions of whether and to what extent perceptual systems operate on stereotyped inputs and perform stereotyped operations. If this is right, then what evidence constitutes support for or against modularity needs to be rethought in light of our conception of basic perceptual strategies.

3 Modularity for Integrativists: The Isotropy Criterion

Even if we are right that an integrative view about perception and modularity *can* be combined, this doesn't show that they should be. Deciding the answer to that question requires a criterion for deciding whether particular perceptual processes/strategies count as modular. Unfortunately, we now want to suggest, the standard criterion for determining whether a perceptual process is modular — pitched in terms of integration across informational boundaries — cannot be straightforwardly employed by an integrativist.

On the standard “boundary-crossing” criterion, employed in virtually all discussions of modularity in the literature, one assesses the putative modularity of perceptual systems in two sequentially ordered steps. First, one draws a functional line around a system in terms of the (unique) feature it purportedly extracts. And second, one asks whether there are genuine informational incursions across that functionally/featurally defined boundary. If so, then the system is non-modular; if not, then it is modular.

Alas, use of this criterion sits quite uneasily with the integrative view of perception. For, on the integrativist view, discovering perceptual strategies means delineating the particular informational interactions that underlie a perceptual process, and describing the computational principles such integrations follow. Given an integrativist conception of perception, it doesn't make sense to attempt to circumscribe some set of features as the genuinely, functionally distinguished, set of inputs, and then regard any use of other inputs as crossing a modular boundary. Rather, the integrationist regards featural boundary-crossing as part-and-parcel of the standard functional operation of perceptual processes. This means that, contrary to the sequence of steps in the standard criterion, integrativists want to draw the functional boundaries after discovering interactions, not before.

On the other hand, if we draw the functional boundaries in a way that makes room for such interactions, as the integrativist prefers, then appeal to the boundary-crossing criterion for modularity threatens to render every process trivially modular. For, the worry goes, there's no amount of interaction that will require giving up modularity — any interactions at all, in any amounts, can instead be counted as taking place inside the (interactive)

modular functional boundaries of the perceptual strategies to which they contribute.¹⁵

These considerations suggest a dilemma for integrativists. Horn one: if we define boundaries prior to detailed investigation of the interactive strategies perception employs, then the discovery of informational interactions will result inevitably in the denial of modularity. Horn two: if we wait until discovering all of the interactions, and then draw the boundaries so as to include relevant informational interactions, then modularity is inevitable. Both horns of this dilemma seem unpalatable, since both trivialize the question of modularity. What's an integrativist to do?

Assuming we do not wish to give up the integrative view, the only way we can see for avoiding the unpalatable options reviewed is to abandon the standard boundary-crossing criterion for modularity. In its place, we propose to return to one of the key ideas Fodor (1983) uses to frame the issue of modularity: the distinction between isotropic and anisotropic processes. The motivating idea here is that what makes modular processes modular, detachable, and in some sense separable from the rest of mentation is that they interface with other aspects of mental processing in a circumscribed number of ways. That is, modular processes are modular just because, and in so far as, there is a delimited range of parameters to which their processing is sensitive. Fodor initially motivates and elaborates the notion of isotropy by the example/analogy of ordinary and scientific rational belief fixation (though he clearly intends to be offering a more widely applicable criterion; cf. Fodor, 1983, 104). What qualifies rational belief fixation as isotropic, Fodor claims, is that there is no circumscribed set of inputs that exhausts the set that is of possible relevance for the updating and revision of any particular belief in the cognitive system, so you have to be prepared to look at information stored anywhere else in the system when fixing belief. Thus, scientific belief fixation is isotropic because "everything that the scientist knows is, in principle, relevant to determining what else he ought to believe. In principle, our botany constrains our astronomy, if only we could think of ways to make them connect" (Fodor, 1983, 105). Generalizing beyond the case of belief fixation, the idea we're proposing is that what qualifies a process as modular is its not being

¹⁵A related objection comes from the worry that the integrativist-friendly individuation procedure will result in an unconstrained, merely disjunctive, characterization of mental processes in terms of literally everything they do — one that will render arbitrary processes trivially modular (Prinz, 2006, 28).

We agree that the procedure in question would trivialize modularity if it depended on utterly unconstrained, merely disjunctive characterizations of mental processes. But our contention is that the integrativist's interaction-including characterizations are *not* unconstrained — that they are, instead, the kinds implicated in the best scientific descriptions of perception. It is worth bearing in mind that whether a particular characterization of a kind is disjunctive depends on the primitive vocabulary. Though *round or red* is disjunctive relative to a vocabulary of primitives such as *round* and *red*, it is also true that *round* is disjunctive relative to a vocabulary of primitives such as *round-and-red* and *round-and-triangular*. As such, just how seriously we take the threat of disjunctivity depends on the vocabulary in terms of which we carry out our theorizing about perception. What the evidence in §1 suggests is that the integrativist's description of perceptual strategies is not disjunctive relative to the kinds enshrined in our best theoretical descriptions of perception.

isotropic — viz., that the range of input parameters to which its processing is sensitive is delimitable.¹⁶

Before we go on to apply the isotropy criterion more widely and draw out its implications, we pause to note that it agrees with standard verdicts about the modularity or otherwise of paradigm cases. Thus, for example, the observation that subjects remain susceptible to the Müller-Lyer illusion even after their rulers or their friends inform them that vision is misinforming them demonstrates that the range of input to which visual length estimation is delimitable (since comparative input from rulers/testimony from friends is ignored). As such, the present criterion classifies the visual processes at issue as anisotropic, so modular. In the other direction, something like analogical reasoning is plausibly not so delimitable, since, in principle, considerations from literally anywhere could be highly relevant to the reliability of analogically transferring conclusions from a first to a second domain. Assuming this is right, the isotropy criterion will classify the processes underlying analogical reasoning as isotropic, hence not modular. It would seem then, that the isotropy criterion at least agrees with standard verdicts about whether these paradigm processes are modular, so is (to that extent) not obviously on the wrong track.¹⁷

A more controversial example to which we can apply our proposed criterion comes from neuroscientific and psychological investigations into action control and prospective memory. These studies are uncovering a system, housed in dorsolateral regions of the prefrontal cortex, that appears to be in charge of representing higher-order relations (e.g., spatio-temporal relations, nesting relations) between different types of perceptual information in the context of action planning and control (Koechlin *et al.*, 2003; Uithol *et al.*, 2012). While much about the representational capacities of this system remains poorly understood/debatable, we want to use it to motivate a general possibility of which the system might turn out to be an instance. Namely, it is at least possible, given what is now known, that this system should turn out to be unrestricted in the range of perceptual parameters to which its processing is sensitive. (Indeed, this possibility makes sense given the possibly infinite set of perceptual objects and features towards which we can direct and organize actions.) If so, then the isotropy criterion would classify this system as non-modular.

We can contrast this case straightaway with the system discussed in §1.2, which integrates contrast, luminance, and object information, underlying the perception of apparent motion. This earlier system seems to operate precisely

¹⁶As Fodor stresses, the fact that a non-delimitable output means that any set of information could in principle be relevant for a particular process does not show, in point of fact, that it is occurrently/actually relevant for that process. Otherwise astronomy would be even more complex than it is now. This will be important for our discussion of cognitive penetrability (§4), where we will emphasize that showing that a system is isotropic — i.e., non-modular — by itself entails nothing about whether particular types of information actually play a role in its processing.

¹⁷Perhaps needless to say, we don't intend to commit to the correctness of the standard verdicts about these cases. Our purpose in appealing to the standard verdicts is only to suggest that the isotropy criterion is not extensionally inadequate in any obvious way.

over these specific inputs, in specific arrangements, and its functioning seems sensitive to these particular sorts of information (and not others). In this case, as in others we've listed, perceptual scientists have usefully attempted to delimit just the types of information that are integrated, and the computational principles via which the integration proceeds. That they are able to do so enables us to describe the operation of the whole in terms of modular perceptual strategies, as per §1.3. This contrasts markedly with the case of the action planning system: here, if we are right that there is no delimitation of the types of information relevant to the processing of the system, an analogous modular description is not in the cards. Whether this is so or not is, of course, an empirical issue — and one we take to be very much up in the air. What we want to emphasize, however, is that the possibility described by characterizing a perceptual system as non-modular is utterly coherent, so not to be dismissed by definition/stipulation.¹⁸

The isotropy framework is meant to open conceptual space for a productive, empirically relevant account of modularity. Many questions, including how to individuate and count modules, how to define relationships between them, etc., remain to be addressed. The question we want to face next, however, is what, if any, implications this unconventional framework for thinking about modularity has for issues about cognitive penetration. This will be the topic of §4.

¹⁸Given this way of thinking about modularity, it is natural to individuate modules by the (by definition, delimited) range of parameters to which their processing is sensitive. Moreover, if what we say about perceptual integration is true (§1), modular systems will overlap — possibly significantly — in their associated ranges of parameters.

It follows from these observations that there are likely to be a large number of modules in perceptual systems. For if module M_1 interacts with or otherwise supplies input to modules M_2 and M_3 , then (some subset of) the parameters figuring in the range associated with M_1 will also figure in the range associated with M_2 and M_3 ; hence it would appear that the range associated with M_1 occurs as a part of not one but many module-individuating feature ranges. That said, we also think that there are explanatory contexts in which it is useful to partition these modules more coarsely, so it's not the case that all of the many modules to which a particular set of parameters contribute are explanatorily relevant in every setting.

A more pressing concern is that, given the extent of perceptual integration, it could turn out that what are intuitively two different modules (say, one canonically associated with vision and another canonically associated with audition) end up being associated with the very same range of parameters; and if so, then it would appear that our view naturally leads to the conflation of intuitively distinct modules. We are not sure that there are genuine cases of this sort. But supposing there are, we suggest that our view can be saved by appealing to some further condition on the individuation of modules. That is, we can say that modules are individuated in terms of both the range of parameters to which their processing is sensitive *and* something else. And there are several different something elses available. For example, one might require that modules be associated with a canonical piece of (neuro-) anatomy, a particular algorithm operating over the inputs, a characteristic form/degree of neural connectivity, etc. As far as we can tell, any of these further conditions would allow us to escape the conflation worry under consideration.

We will ignore these niceties in what follows. (Thanks to Matthew Fulkerson for discussion of these matters.)

4 Higher perception, modularity, and cognitive penetrability

One important ramification of thinking about the modular/non-modular distinction in terms of isotropy, as we propose, is that that distinction comes apart from the cognitive/non-cognitive distinction. This is relevant because virtually all contemporary discussions of cognitive penetration (by both proponents and opponents) of which we're aware have turned on supposing that these distinctions coincide, so that one can move freely between conclusions about modularity and conclusions about cognitive penetration.

The literature about cognitive penetration has to date largely concerned the assessment of whether there is or is not crossing of a modular boundary (viz., a boundary demarcating perception) by ostensibly cognitive factors in a handful of cases. For example, Macpherson (2012) appeals to cases of interactions between shape and color — specifically the result that heart shapes are perceived as redder than non-heart shapes of similar size — to argue for cognitive penetration of color experience. And Stokes (2012) argues from the effect of increased size judgments for coins versus similarly shaped disks that size perception is mediated by value judgments. In each of their cases, subjects are asked to match the stimulus to a comparator — the background in the color case, or another disk in the shape case — that can be adjusted manually by the subject until the test stimulus and comparator look the same along the relevant dimension. Crucially, they take the effects in question to turn on categorical judgments that, according to them, can only be regarded as having been carried out by cognitive, rather than perceptual, processes:

While it may be plausible to think that the visual system is sensitive to heart shapes and autonomously responds to that shape making it seem more red than it is, it is far less plausible to think that the visual system can classify certain objects as being apples and autonomously alter the apparent colour of such objects. The reason is that classifying something as an apple doesn't depend just on simple features, such as shape, that the visual system is responsive to, for many things have the shape of apples, such as cherries and other fruits. And while the visual system can respond to shape, it surely can't classify an object as an apple, as it is not sensitive to the features that are required to do so. It is surely only the cognitive system that is capable of doing such classificatory work (Macpherson, 2012, 47–48).

On the other side, those who deny the existence of cognitive penetration typically attempt to undermine alleged cases of penetration by arguing that the latter do not involve the crossing of a modular boundary — e.g., by arguing that, in such cases of apparent unencapsulation, the relevant effects of cognition are due either to attentional processes that operate before perceptual modules are engaged, or to post-perceptual judgments that operate after

perceptual modules have finished (e.g. Pylyshyn, 1999).¹⁹ Whatever one's view about the success of these defenses, the point to notice is that these theorists are accepting the inference from a lack of modularity to cognitive penetration: they are agreeing with their opponents that *if there were* a crossing of a modular boundary, then that *would be* sufficient to demonstrate that perception is penetrated by cognition. Indeed, it is precisely because they accept this inferential pattern, that their rejection of cognitive penetration depends on finding ways to deny that a modular boundary is crossed in the cases under discussion.

Thus, while these theorists disagree about whether there are genuine cases of non-modularity, all parties to the dispute accept the goodness of the inference from non-modularity to cognitive penetration, and so confine their disagreement to the issue of whether certain cases qualify. In what follows we will suggest that the crucial inference underlying the dispute — and acceded by theorists on both sides — looks significantly less attractive once we think about modularity in the terms we have been urging.

Before we come to this, however, we want to register a more internal concern. Namely, we think there are reasons to doubt that the classificatory status of the output of a representational process, invoked explicitly by Macpherson in the passage quoted above, is a good criterion for that state's being cognitive (/cognitively penetrated).²⁰ To see why, consider investigations into the phenomenon of perceived "chasing" (Gao *et al.*, 2009, 2010; Gao and Scholl, 2011) which have shown that certain patterns of movement between objects, even in the presence of many distracters, reliably elicit the representation that a second object is chasing a first. These investigations reveal that verdicts about chasing, no less than verdicts about applehood, seem to depend on more than just the simple features to which vision is responsive. Thus, for instance, if objects are not "facing" their direction of movement (e.g., if a triangle's

¹⁹For what it is worth, Macpherson and Stokes address and argue against the possibility of defending impenetrability by these strategies in the cases they discuss. They argue against the attentional/pre-perceptual explanation on the grounds that, in the experiments cited, subjects have as much time as they need to adjust the comparator, hence have ample time for a wide variety of shifts in attention. They argue against the judgment/post-perceptual explanation on the grounds that this story would result in a deep inconsistency between the perceptual and judgmental interpretations if the affected versus unaffected stimuli were perceived as the same. We take no official stand on the success of these arguments.

²⁰The version of the challenge to which we respond in the main text, taken from Macpherson's description quoted above, revolves around the idea that the subject's response is specifically *classificatory*. But there is a distinct, related formulation of the challenge grounded in the idea that the subject's response is *categorical*.

Besides wanting to follow Macpherson's formulation of the worry, we prefer to carry out the discussion in terms of classification rather than categoricity because the latter tends to carry stronger conceptual/cognitive associations, and we are precisely attempting to leave it up for grabs whether and to what extent concepts/cognition play a role in generating the responses at issue. That said, note that our response to the case, which involves the example of chasing (see below), plausibly extends to the version of the challenge formulated in terms of categoricity as well. For, though one might initially have described subjects' representations of whether one object is chasing another as both classificatory and categorical, our suggestion is that, on the evidence, these representations may be capable of being understood in non-cognitive terms.

lowest-degree angle is oriented away from the direction in which it moves), this inhibits the classification of the case as an instance of chasing. More generally speaking, it appears that the effect only occurs for figures whose shapes, as well as motions, stand in particular configural relations to one other. But despite being apparently classificatory in the intended sense, it is by no means obvious that representations of chasing (/the processes underlying that sort of representation) are cognitive. Indeed, while it is too early to endorse this conclusion with full confidence, initial research suggests strongly that the representation of chasing can be accounted for in terms of representations of shapes and motions (and configural relations among these) that are ordinarily counted as wholly perceptual, as opposed to beliefs or other canonically cognitive states. (For what it is worth, Gao and Scholl (2011, 669) themselves advocate understanding the representation of chasing in perceptual rather than cognitive terms.)

If we are right that representations of chasing might be examples of classificatory representations that are not themselves cognitive or cognitively mediated, then this would defeat the claim that any perceptual states sensitive to classificatory effects are *ipso facto* evidence of informational exchange across a modular boundary. In turn, this would mean that, even granting the boundary crossing conception of modularity we rejected (§2), cases of perceptual states sensitive to classificatory effects can't be assumed without further argument to represent modularity violations, or, therefore, grounds for concluding that cognition penetrates perception.

Having registered this internal concern, we now want to raise a further, and less internal, point suggested by the case of representations of chasing: the processing underlying such classificatory/high level representations might well turn out to be anisotropic. For, while, once again, it is still early days, the initial evidence suggests strongly that the inputs for the processes leading to classificatory verdicts about chasing are delimited and specifiable. Therefore, on the isotropy construal urged in §2, these processes will plausibly turn out to be modular despite producing output representations that are no less classificatory than those considered by Macpherson and Stokes. Nor is chasing unique in being both plausibly anisotropic and ostensibly classificatory/high-level: investigators have proposed that the perception of certain social cues (Langton *et al.*, 2000) and biological motion (Johansson, 1973; Troje and Westhoff, 2006) may admit of similarly modular interpretations (though, once again, we don't mean to hold our conceptual point hostage to these verdicts about the particular phenomena). What this suggests, then, is that the classificatory nature of a representation is not a reliable guide to its being either cognitively penetrated or modular.

Indeed, these considerations put us in a position to make the even stronger and more direct suggestion that (whatever one thinks about classificatory representations) at least some stereotypically cognitive processes might rely on a significantly delimitable range of input parameters — hence, again, assuming our revisionary construal, might turn out to be modular. We suggest that simple arithmetic, for example, might be a case of this sort: it might well

turn out that mental arithmetic draws on a delimited suite of psychological capacities, though these capacities and the associated contents on which they operate will plausibly be classified as conceptual rather perceptual (for an overview of research in this area, see Ashcraft, 1992). This gives reason for thinking that, on an isotropy-based reconstrual of modularity, the standard inference from modularity to cognitive impenetrability breaks down.

Nor, given our preferred construal of modularity, is there reason to expect that non-modular (viz., isotropic) perceptual processing must be sensitive to the cognitive system. Thus, as we suggested in the case of prospective memory and action planning (§2), it is least possible (and perhaps actual) that systems involved in the organization of perceptual information might draw on a non-delimitable range of inputs despite not drawing on inputs from cognition.²¹ This conclusion appears to comport well with recent work on perceptual category learning, which suggests that the processing underlying perceptual category learning becomes increasingly sensitive to integrative (“contextual”) cues — cues involving not just the presence of isolated parts of the stimulus, but the configural relations holding between parts — as subjects become more expert (Gauthier and Tarr, 2002). Given the large (perhaps unlimited) number of objects about which we can form perceptual expertise, it is likely that this system won’t have delimitable inputs. However as we saw with prospective memory, this point offers no reason to believe that the processing of configural cues underlying the resulting categorical representations must reflect input from cognition in particular. Hence there is no good inference here from the isotropy of the system to its having been penetrated by cognition. That is, there is no good reason to suppose, simply from the fact that such processes operate over a non-delimitable range of inputs, and that the outcome of learning over these inputs implements categorical effects, that they are sensitive to *cognitive* representations in particular.²²

Now, the point we’re urging requires some care: for, while what we have said shows that there is one understanding of cognitive penetration on which

²¹One way in which one might wish to challenge our description of this system is by claiming that it is in fact cognitively penetrated, so not an instance of non-modularity without cognitive penetration after all (thus, for example, cf. Wu (2012) for an analogous argument that the system for visual spatial constancy is cognitively penetrated by intentions). While we cannot hope to resolve the issue about prospective memory in particular (or about Wu’s case of visual spatial constancy) here, we stand by our claim in the main text: construed our way, the non-modularity of a system is not by itself a guarantee of its cognitive penetration.

²²Here we are assuming that there is a criterion of the cognitive — a standard that would settle whether a class of representations that affects some output is cognitive or not. For example, on the standard assumption that explicitly believed content is cognitively represented, one might test for the influence of specifically cognitive representations by determining whether perceptual processes are sensitive to (some degree of) relevant change in explicit belief. Thus, if the size or color judgments in the cases discussed could be mediated by previous belief change (e.g., being told that the coin was a counterfeit), then this would support the view that the interactions in question are cognitive. These are only suggestions: we take no official stand here about the nature of the cognitive (or about the appropriate experimental tests for cognitive status). Rather, our central point is that the question of whether a class of representations is cognitive should proceed independently of the question of whether a particular process affected by members of that class is modular or not.

the inference from isotropy to penetration fails, there is a neighboring reading of cognitive penetration on which that inference withstands the considerations we have brought to bear. Specifically, we have suggested that an isotropic process can fail to draw on cognitive input, so that such processes can fail to be cognitively penetrated as a matter of actual, historical fact. On the other hand, the isotropy of such processes means that we cannot foreclose, *ex ante*, the possibility that they *might* draw on — and that their operation would be affected by drawing on — any particular range of input, including, *inter alia*, cognition. In slogan form, even if our considerations suggest that isotropic processes might fail to be (actually) cognitively *penetrated*, they don't show that such processes might fail to be cognitively *penetrable*.

Even with this concession, however, we take our considerations to be relevant to the kind of reasoning on display in many recent discussions (e.g. Macpherson, 2012; Stokes, 2012), which attempts to infer from the absence of modularity to the conclusion that perception is actually cognitively penetrated (and thence, presumably, to the corollary conclusion that perception is cognitively penetrable, by application of a trivial modal axiom).²³

To be clear, our contention is not that there is no penetration by cognition going on in these cases; for we take the ultimate extent, type, and import of the relations between perception and cognition to be very much an open question. Rather, our point is that once modularity is construed in terms of isotropy, the proffered conclusion does not follow from the premises provided: the isotropy of a process, *per se*, is no guarantee of its (actual) penetration by cognition.

5 Conclusion

We have argued that standard discussions of cognitive penetrability, perception, and modularity rest on an insufficient appreciation of the extent and role of integration within perception. But, unlike some who would take such considerations as reason for rejecting modularity altogether, we take them instead to motivate a more refined, isotropy-based conception of modularity, as opposed to the usual boundary-crossing criterion.

Crucially, however, adopting the newer criterion for modularity means recognizing that the modularity/non-modularity distinction comes apart from the cognitive/non-cognitive distinction, and that they together partition mental processes into four kinds rather than two. It allows for processes that are non-modular and cognitively penetrated (e.g., rational belief fixation); processes that are non-modular and not cognitively penetrated (e.g., prospective memory); processes that are modular and cognitively penetrated (e.g., mental arithmetic); and processes that are modular and not cognitively penetrated (the representation of chasing). We think modifying our understanding of

²³In fairness, Macpherson and Stokes explicitly argue for cognitive penetration from the premise that informational encapsulation fails, not from the premise that modularity fails. But since these authors interdefine the notions of modularity and informational encapsulation, this is a distinction without a difference.

modularity in this way, and accepting the resulting taxonomic revisions, carries significant benefits. It offers gains in clarity, allows us to recognize the importance of integration within perception without requiring us to give up useful notions of modularity, and leads to a more expansive view of the possible ways in which mental/perceptual states can interact.

Our view leaves open important questions about how modular and non-modular systems interact within and across the perceptual/cognitive divide, just how that divide should be construed once the integrative picture is accepted, and what conditions distinguish modular cognitive systems from modular perceptual systems. But we believe that it represents a promising framework for pursuing these questions, and, more generally, for thinking about modularity, cognitive penetration, and mental architecture.²⁴

References

- Anstis, S. M. (2003a). Levels of motion perception. In L. Harris and M. Jenkin, editors, *Levels of perception*, pages 75–99. Springer, New York.
- Anstis, S. M. (2003b). Moving objects appear to slow down at low contrasts. *Neural Networks*, **16**, 933–938.
- Anstis, S. M. (2004). Factors affecting footsteps: contrast can change the apparent speed, amplitude and direction of motion. *Vision Research*, **44**, 2171–2178.
- Ashcraft, M. H. (1992). Cognitive arithmetic: A review of data and theory. *Cognition*, **44**, 75–106.
- Bach y Rita, P., Collins, C. C., Saunders, F., White, B., and Scadden, L. (1969). Vision substitution by tactile image projection. *Nature*, **221**, 963–964.
- Barrett, H. C. and Kurzban, R. (2006). Modularity in cognition: Framing the debate. *Psychological Review*, **113**(3), 628–647.
- Broackes, J. (2009). What do the colour-blind see? In J. Cohen and M. Matthen, editors, *Color Ontology and Color Science*. MIT Press, Cambridge, Massachusetts.
- Buller, D. J. (2005). *Adapting Minds: Evolutionary Psychology and the Persistent Quest for Human Nature*. MIT Press, Cambridge, Massachusetts.
- Buller, D. J. and Hardcastle, V. G. (2000). Evolutionary psychology, meet developmental neurobiology: Against promiscuous modularity. *Brain and Mind*, **1**(3), 307–325.

²⁴Thanks to Matthew Fulkerson, Brian Keane, Casey O’Callaghan, and to the editors of this volume, for helpful discussions of this material that significantly improved the paper.

- Burnston, D. and Cohen, J. (2013). Perception of features and perception of objects. *The Croatian Journal of Philosophy*. in press.
- Carruthers, P. (2006). *The architecture of the mind: Massive modularity and the flexibility of thought*. Oxford, Oxford University Press.
- Cohen, J. (2012). Perceptual constancy. In M. Matthen, editor, *Oxford Handbook of Philosophy of Perception*. Oxford University Press, Oxford. forthcoming.
- Coltheart, M. (1999). Modularity and cognition. *Trends in Cognitive Sciences*, **3**(3), 115–120.
- Field, G. D., Greschner, M., Gauthier, J. L., Rangel, C., Shlens, J., Sher, A., Marshak, D. W., Litke, A. M., and Chichilnisky, E. J. (2009). High-sensitivity rod photoreceptor input to the blue-yellow color opponent pathway in macaque retina. *Nature Neuroscience*, **12**, 1159–1164.
- Fodor, J. A. (1983). *The Modularity of Mind*. MIT Press, Cambridge, Massachusetts.
- Fodor, J. A. (2001). *The mind doesn't work that way: The scope and limits of computational psychology*. MIT Press, Cambridge, Massachusetts.
- Gao, T. and Scholl, B. J. (2011). Chasing vs. stalking: Interrupting the perception of animacy. *Journal of Experimental Psychology: Human Perception and Performance*, **37**, 669–684.
- Gao, T., Newman, G. E., and Scholl, B. J. (2009). The psychophysics of chasing: A case study in the perception of animacy. *Cognitive Psychology*, **59**, 154–179.
- Gao, T., McCarthy, G., and Scholl, B. J. (2010). The wolfpack effect: Perception of animacy irresistibly influences interactive behavior. *Psychological Science*, **21**, 1845–1853.
- Gauthier, I. and Tarr, M. J. (2002). Unraveling mechanisms for expert object recognition: Bridging brain activity and behavior. *Journal of Experimental Psychology: Human Perception and Performance*, **28**(2), 431–446.
- Geldard, F. A. and Sherrick, C. E. (1972). The cutaneous "rabbit": A perceptual illusion. *Science*, **178**(4057), 178–179.
- James, T. W., Humphrey, G. K., Gati, J. S., Servos, P., Menon, R. S., and Goodale, M. A. (2002). Haptic study of three-dimensional objects activates extrastriate visual areas. *Neuropsychologia*, **40**(10), 1706–1714.
- Johansson, G. (1973). Visual perception of biological motion and a model for its analysis. *Perception & Psychophysics*, **14**(2), 201–211.
- Kayser, C. and Logothetis, N. K. (2007). Do early sensory cortices integrate cross-modal information? *Brain Structure and Function*, **212**(2), 121–132.

- Koechlin, E., Ody, C., and Kouneiher, F. (2003). The architecture of cognitive control in the human prefrontal cortex. *Science*, **203**(5648), 1181–1185.
- Kohly, R. P. and Regan, D. (2002). Fast long-range interactions in the early processing of luminance-defined form. *Vision Research*, **42**(1), 49–63.
- Langton, S. R. H., Watt, R. J., and Bruce, V. (2000). Do the eyes have it? cues to the direction of social attention. *Trends in Cognitive Sciences*, **4**(2), 50–59.
- Lotto, R. B. and Purves, D. (2001). An empirical explanation of the Chubb illusion. *Journal of Cognitive Neuroscience*, **13**(5), 547–555.
- MacKay, D. (1958). Perceptual stability of a stroboscopically lit visual field containing self-luminous objects. *Nature*, **181**, 507–508.
- Macpherson, F. (2012). Cognitive penetration of colour experience: Rethinking the issue in light of an indirect mechanism. *Philosophy and Phenomenological Research*, **84**(1), 24–62.
- Mesulam, M. M. (1998). From sensation to cognition. *Brain: A journal of neurology*, **121**(6), 1013–1052.
- Meyer, K. and Damasio, A. (2009). Convergence and divergence in a neural architecture for recognition and memory. *Trends in Neurosciences*, **32**(7), 376–382.
- Nijhawan, R. (1994). Motion extrapolation in catching. *Nature*, **370**, 256–257.
- O’Callaghan, C. (2013). Perception and multimodality. In E. Margolis, R. Samuels, and S. Stich, editors, *Oxford Handbook to Philosophy and Cognitive Science*, chapter four. Oxford University Press, Oxford.
- Pick, Jr., H. L., Warren, D. H., and Hay, J. C. (1969). Sensory conflict in judgements of spatial direction. *Attention, Perception, & Psychophysics*, **6**(4), 203–205.
- Prinz, J. J. (2006). Is the mind really modular? In R. Stainton, editor, *Contemporary Debates in Cognitive Science*, pages 22–36. Blackwell, Oxford.
- Purves, D., Shimpi, A., and Lotto, R. B. (1999). An empirical explanation of the Cornsweet effect. *The Journal of Neuroscience*, **19**(19), 8542–8551.
- Pylyshyn, Z. W. (1999). Is vision continuous with cognition? the case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, **22**(3), 341–423.
- Samuels, R. (2006). Is the human mind massively modular? In R. Stainton, editor, *Contemporary debates in cognitive science*, pages 37–56. Blackwell, Cambridge, Massachusetts.

- Shams, L., Wozny, D. R., Kim, R., and Seitz, A. (2011). Influences of multisensory experience on subsequent unisensory processing. *Frontiers in Psychology*, **2**(264).
- Shepard, R. N. (2001). Perceptual-cognitive universals as reactions of the world. *Behavioral and Brain Sciences*, **24**, 581–601.
- Sperber, D. (2005). Modularity and relevance. In P. Carruthers, S. Laurence, and S. Stich, editors, *The innate mind: Structure and contents: Volume 1*, pages 53–68. Oxford University Press, Oxford.
- Stein, B. E. and Stanford, T. R. (2008). Multisensory integration: current issues from the perspective of the single neuron. *Nature Reviews Neuroscience*, **9**, 255–266.
- Stokes, D. (2012). Perceiving and desiring: A new look at the cognitive penetrability of experience. *Philosophical Studies*, **158**(3), 479–92.
- Troje, N. F. and Westhoff, C. (2006). The inversion effect in biological motion perception: Evidence for a “life detector”? *Current Biology*, **16**, 821–824.
- Uithol, S., Burnston, D., and Haselager, P. (2012). Will intentions be found in the brain? under review.
- Ullman, S. (1979). *The Interpretation of Visual Motion*. MIT Press, Cambridge, Massachusetts.
- Vroomen, J. and de Gelder, B. (2000). Sound enhances visual perception: cross-modal effects of auditory organization on vision. *Journal of Experimental Psychology: Human Perception and Performance*, **26**(5), 1583–1590.
- Vroomen, J. and de Gelder, B. (2004). Ventriloquism and the freezing phenomenon. In G. A. Calvert, C. Spence, and B. E. Stein, editors, *The Handbook of Multisensory Processes*, pages 141–150. MIT Press, Cambridge, MA.
- Wu, W. (2012). Visual spatial constancy and modularity: Does intention penetrate vision. *Philosophical Studies*.
- Young, M. P., Hilgetag, C. C., and Scannell, J. W. (2000). On imputing function to structure from the behavioural effects of brain lesions. *Philosophical transactions of the Royal Society of London Series B, Biological sciences*, **355**(1393), 147–161.
- Yuille, A. and Kersten, D. (2006). Vision as Bayesian inference: analysis by synthesis? *Trends in Cognitive Sciences*, **10**, 302–308.