



Cognitive penetration and informational encapsulation: Have we been failing the module?

Sam Clarke¹

Accepted: 28 September 2020
© Springer Nature B.V. 2020

Abstract Jerry Fodor deemed informational encapsulation ‘the essence’ of a system’s modularity and argued that human perceptual processing comprises modular systems, thus construed. Nowadays, his conclusion is widely challenged. Often, this is because experimental work is seen to somehow demonstrate the cognitive penetrability of perceptual processing, where this is assumed to conflict with the informational encapsulation of perceptual systems. Here, I deny the conflict, proposing that cognitive penetration need not have any straightforward bearing on (a) the conjecture that perceptual processing is composed of nothing but informationally encapsulated modules, (b) the conjecture that each and every perceptual computation is performed by an informationally encapsulated module, and (c) the consequences perceptual encapsulation was traditionally expected to have for a perception-cognition border, the epistemology of perception and cognitive science. With these points in view, I propose that particularly plausible cases of cognitive penetration would actually seem to evince the encapsulation of perceptual systems rather than refute/problematize this conjecture.

Keywords Modularity · Cognitive penetration · Informational encapsulation · Mental imagery · Fodor (Jerry)

This paper was much improved by the constructive and critical comments I received from Greyson Abid, Jake Beck, Steve Butterfill, Laurenz Casser, Martin Davies, Azenet Lopez, Fiona Macpherson, Mike Martin, Ian Phillips, Jake Quilty-Dunn, Nick Shea, an anonymous reviewer, and audiences in Athens, Austin, Glasgow, Oxford, and Toronto. The research was supported by VISTA Research Fellowship at York University.

✉ Sam Clarke
sclarke@yorku.ca

¹ Department of Philosophy and Centre for Vision Research, York University, 247 Lippincott Street, Toronto, ON, Canada

1 Introduction

Is perception modular? Jerry Fodor thought so. In a seminal (1983) contribution he clarified the suggestion, deeming *informational encapsulation* ‘the essence’ of a system’s modularity (p. 71). He then argued that (as a matter of empirical fact) perceptual processing comprises modular systems, thus construed (p. 101).

This refined earlier appeals to the functional specialisation of these systems (e.g. Marr 1982, p. 325), was expected to mark a perception-cognition border (Fodor 1983, 1985), and was expected to have important implications for the epistemology (Fodor 1983, 1984) and computational tractability of perception (Fodor 1983, 2000). The thing is: Fodor’s conclusion is widely challenged. Standardly, this is because experimental work is seen to somehow demonstrate the cognitive penetrability of perceptual processing. Since cognitive penetration is widely assumed to be in tension with the informational encapsulation of perceptual systems, this has prompted many to abandon the modularity of perceptual processing altogether (Churchland 1988; Lupyán 2015; Wu 2013) or to propose watered-down notions of modularity that take little or no stand on informational encapsulation (Coltheart 1999; Lyons 2015; Ogilvie and Carruthers 2015). For some, this amounts to a change of subject (cf. Wilson 2008). But even when theorists defend the encapsulation of perceptual systems, they (Fodor included) tend to do so by rejecting the suggestion that cognitive penetration actually occurs (e.g. Firestone and Scholl 2016; Fodor 1983, 1988; Pylyshyn 1999; but see Quilty-Dunn 2020). So, while their conclusions diverge, theorists of all stripes proceed as if cognitive penetrability were problematic for, and perhaps incompatible with, the claim that perceptual processing is composed of encapsulated modules.

The purpose of this paper is to show that this is a mistake. I defend the claim that human perception is modular where this involves (a) human perceptual processing composing of nothing but encapsulated modules and (b) the full-blown encapsulation of each and every computation taking place in human perception. In spite of this, I remain neutral on the state of the evidence for cognitive penetration. I argue that even if cognitive penetration is pervasive in the way some claim, it is a mistake to assume that this must threaten (a) and/or (b). Indeed, I propose that it need not even threaten the interesting consequences perceptual modularity/encapsulation was expected to have for a perception-cognition border, the epistemology of perception, and the computational tractability of perceptual processing. With these suggestions made (henceforth, my *conceptual point*), I argue for a bold claim: particularly plausible cases of cognitive penetration would actually seem to evince (a) and (b), rather than undermine or refute these claims (I call this my *empirical point*).

I proceed as follows: In §2, I introduce the claim that perception is cognitively penetrable (§2.1) and Fodor’s claim that perceptual processing is composed of informationally encapsulated modules (§2.2). I then note one way in which perceptual processing could comprise nothing but encapsulated modules, despite widespread cognitive penetration (§2.3). To this end, I consider the possibility that perception is underwritten by a hierarchy of modules, wherein outputs of lower-level modules provide inputs to modules located at higher levels of processing. On a

picture of this sort, cognitive penetration could occur at the joints between modules in the hierarchy, leaving each module completely encapsulated from central cognition and the rest of the mind. In addition, I argue that the perceptual modifications themselves (resulting from cognitive penetration occurring at the joints between modules in the hierarchy) could avoid amounting to perceptual computations, the upshot being that perception could comprise nothing but encapsulated modules, performing fully encapsulated computations, despite pervasive cognitive penetration. After arguing that this does not change the subject—and that it need not undermine the interesting consequences perceptual modularity was originally expected to have for a perception-cognition border, the computational tractability of perceptual processing, and the epistemology of perception (§2.4)—I turn to my empirical point: identifying a particularly plausible case of cognitive penetration (penetration mediated by mental imagery) and arguing that (if actual) this would evince a modular architecture of the aforementioned sort, rather than refute or undermine its postulation (§3.1). I close by considering how opponents might put pressure on the proposal (§3.2).

2 The conceptual point

I begin by introducing my terms.

2.1 Perception as cognitively penetrable

It is largely uncontroversial that perception influences cognition. For instance, when I perceive a red blob in front of me, as such, this disposes me to believe (i.e. cognise) that there is a red blob in front of me (for discussion, see Johnston 1992, p. 222; Smith 2001, p. 291; Tye 1995, pp. 143–4). To a first approximation, the claim ‘perception is cognitively penetrable’ amounts to the more surprising suggestion that information flows in the opposite direction: that one’s cognitive states (what one believes, desires, intends, etc.) can systematically influence one’s perception or perceptual processing.

This requires qualification. The existence of cognitive penetration is a contentious issue, and some claim its existence would call for a revolution in the way we study perception, requiring researchers to begin carefully considering subjects’ beliefs, desires and expectations when investigating their perceptual psychology (Firestone and Scholl 2016). There are, however, many mundane ways that cognition influences perception which are unsurprising and fail to have radical consequences of this sort. Thus, in order to specify the phenomenon at issue, theorists typically hold that ‘cognitive penetration’ (as they are interested in it) would involve one’s cognitive states exerting a *direct* and *semantically coherent* effect on their perceptual processing. In any case, it is cognitive penetration of this sort which is typically presented as problematic for the modularity of perceptual systems.

For one’s cognitive states to exert a *direct* effect on one’s perceptual processing, each step in the causal process would need to occur internally to the organism and at

a psychological level of description (Stokes 2012). This rules out trivial and uninteresting cognitive influences of the following sort: Mihir intentionally moves his head, and this results in changes to the content/character of his visual processing. Here, Mihir's perceptual processing has been affected by one or more of his cognitive states (most obviously, an intention). In spite of this, it would be rash to assume that his cognitive states have thereby 'penetrated' his perceptual processing in any theoretically interesting sense of the term. This is because it remains plausible that his cognitive states have simply guided an action that has led to different objects stimulating his retinae and, thus, getting interpreted by perceptual systems (see Firestone and Scholl 2016, pp. 3–4). As such, it is plausible that important steps in the causal process are non-psychological and occurring outside of the organism (e.g. Mihir's perceptual processing is being influenced by non-psychological effects at his sensory organs). In this way, Mihir's cognitive states would seem to be influencing his perceptual processing. But, *prima facie*, they are only doing so *indirectly*.

Indirect effects of this sort are widespread, uncontroversial and carefully controlled for by perception scientists. For this reason, they are not at issue in debates concerning cognitive penetration. They can, however, be contrasted with conceivable effects of the following sort: Aga is looking at a coin and, by believing it to be valuable, comes to perceive the coin as bigger than she would have done otherwise (c.f. Bruner and Goodman 1947). Here, Aga's belief (a cognitive state) may seem to have affected her visual processing independently of what Aga actually perceives. Consequently, it may seem unlikely that the perceptual alteration could be attributed to a simple change in sensory input (for discussion, see Pylyshyn 1999, and commentaries). Instead, it may seem that Aga's belief is influencing the way(s) in which these inputs are being interpreted by her visual systems, themselves. In this way, Aga's cognitive states could seem to be 'directly' influencing her perceptual processing, in that the causal influence would be occurring internally to the organism (Aga) and at a psychological level of description. Since effects of this sort are non-obvious and would be hard to control for in psychological experiments, their existence may have important consequences for how we study and think about perception (Firestone and Scholl 2016).

The claim that cognitive penetration requires a *semantically coherent* effect (Pylyshyn 1999, p. 343) concerns the suggestion that the perceptual effect must also bear a non-accidental, content respecting relation to the cognitive state that brought it about (Macpherson 2015). This is intended to preclude cases of the following sort: Betty believes that she will fail her exam and this causes her to have a migraine where this migraine causes her visual system to represent flashing lights in the periphery of her visual field (Macpherson 2012, p. 26). Here, a belief could be said to have directly affected Betty's perceptual processing in the above way (we might allow that the relevant stages of the causal process—however numerous—all occur within Betty's cranium and at a psychological level of description). In spite of this, the criterion of semantic coherence would seem to prevent this from qualifying as genuine cognitive penetration. Why? Because there would not seem to be a relevant, rational, or logical relation between the content of the offending cognitive state (an impending exam) and the modified content of perception (flashing lights).

Admittedly, our assessment of these cases may rely on intuition. Thus, the above point may simply amount to the claim that, *intuitively*, exams bear no semantic relevance, or logical connection, to flashing lights (a point some find unsatisfactory—Stokes 2013; but see Macpherson 2015). Regardless, the example can be contrasted with one in which the relevance seems fairly obvious. For instance, we can conceive of a situation in which Neil’s belief that ‘heart shapes tend to be coloured red’ exerts a direct influence on his visual systems’ interpretation of an orange heart shape, causing it to represent the shape as *red* or *redder* than it would have done otherwise (see Delk and Fillenbaum 1965). Here, the logical connection between the cognitive state and the perceptual modification would seem relatively clear. Consequently, it is effects of this sort that proponents of cognitive penetration purport to identify (Lupyan 2015; Mole 2015; Wu 2013) and proponents of perception’s cognitive *impenetrability* deny exist (Firestone and Scholl 2016; Gross et al. 2013; Gross 2017; Pylyshyn 1999). So, while there is debate over the existence of cognitive penetration, it is often agreed that cognitive penetration would occur if (and perhaps only if) an organism’s cognitive states exerted a direct and semantically coherent effect on their perceptual processing. In any case, it is cognitive penetration of this sort that has typically been seen to problematize the informational encapsulation of perceptual systems. Unless otherwise noted, this is how I will use the term.¹

2.2 Perceptual systems as informationally encapsulated

Informational encapsulation is a quite different matter. It pertains to a *restriction* on the information a system has access to in its computations. So, while ‘cognitive penetration’ refers to a type of *freedom* in the flow of information between cognition and perception (specifically, a freedom for the former to affect the latter in direct and semantically coherent ways), a system A (e.g. a visual system) is encapsulated from a system B (e.g. central cognition) insofar as there is a particular type of restriction on system A’s access to the information accessible to system B.

To appreciate the nature of this restriction, and Fodor’s conjecture that perceptual processing is composed of encapsulated modules, some background is helpful.

First, when Fodor introduced the term ‘informational encapsulation’ he was interested in characterising restrictions on the information accessible to relatively fine-grained perceptual sub-systems. For while theorists often speak as if each perceptual modality (or even perception as a whole) was supposed to constitute its own encapsulated module, this was never Fodor’s “intended doctrine” (1983, p. 47). His proposal was that “within (and quite possibly across) the traditional

¹ Cognitively penetrability’ is sometimes seen to concern the cognitive penetrability of *perceptual experience* (see Macpherson, 2012, p. 27). There, the question is whether the ‘what it’s like’ (Nagel 1974) of perceptual *experience* is affected by our cognitive states in relevant ways. This is not how I am understanding matters here. Since it is unclear how the modularity of perception relates to perceptual experience—some modular outputs may be unconscious (Mandelbaum 2018; Quilty-Dunn 2019), and it is possible that all are (Mandelbaum 2019)—the important question (for my purposes) is whether an organism’s cognitive states might directly affect their perceptual processing in a semantically coherent way, irrespective of whether this results in modifications of experience (Pylyshyn 1999).

modes, there are highly specialised computational mechanisms” where the “specialization of these mechanisms consists in restrictions on the information they can access” (*ibid.*). For Fodor, it was these specialised mechanisms that were the modular systems of human perception (p. 48) and whose encapsulation he was trying to capture. To say that perceptual “input” systems are encapsulated and modular was, thus, to say that perceptual processing is (entirely) composed of such mechanisms, where each mechanism is relevantly restricted in its access to information.²

Second, Fodor was working on the assumption that sensory inputs underdetermine the eventual analyses of perception (Fodor and Pylyshyn 1981; cf. Helmholtz 1876). While this has not gone unchallenged (Gibson 1979), it remains a basic tenet of mainstream perception science (see Palmer 1999; Rescorla 2015). And, as Fodor was keen to note, it implies that the aforementioned perceptual mechanisms engage in a form of “non-demonstrative inference” (1983, p. 68). Why? Because if sensory inputs systematically underdetermine perceptual outputs, the non-accidental reliability, accuracy and utility of these outputs would seem to depend on the relevant perceptual systems making some kind of informed prediction about the causes of their inputs. In this respect, each of the perceptual mechanisms in question was deemed a “computational mechanism that projects and confirms a certain class of hypotheses on the basis of a certain body of data” (*ibid.*).

This is now common ground among Fodor and his opponents (at least those with whom we are engaging). But notice: the ‘body of data’ that informs these perceptual analyses could always be more or less unbounded. In principle, it could include all things known or believed by the organism and its other psychological systems (see Bruner and Postman 1949; Halle and Stevens 1962). Thus, it is conceivable that a visual system would infer shape from shading on the basis of (e.g.) the organism’s explicit belief that ‘light tends to come from above’ (Brown and Friston 2012; Clark 2015, pp. 85–86), or even their beliefs about weather and astrology (Fodor 1983, p. 64).

Fodor conjectured otherwise. He proposed that the ‘body of data’ that is accessible to each of the specialised mechanisms in question, and (thus) used by these when interpreting their inputs, happens to “include, in the general case, considerably less than the organism may know” (1983, p. 69). In each case, it is restricted to an architecturally prescribed ‘proprietary database’ that is dedicated to that mechanism’s operations and (typically) that mechanism’s operations alone (Fodor 1985, p. 3). In this respect, a system A qualifies as ‘informationally encapsulated’ from a system B, insofar as the information it has access to is limited to that which is located inside its proprietary database, and to the extent that this proprietary database fails to include information that is accessible to system B. According to Fodor, perceptual processing is notable in that it is (entirely)

² This is not to deny that perception as a whole may still qualify as a unified *system* by some legitimate criterion, even if it doesn’t count as a single system by virtue of its constituting a single module, with similar points applying to individual sense modalities. The present treatment takes no stand on these issues. To say that ‘perceptual systems are modular’ I simply mean to maintain that perceptual processing is entirely composed of modular systems.

composed of modules that are encapsulated from the system(s) of central cognition in this respect.³

The Muller-Lyer illusion is often used to illustrate this suggestion (see Fig. 1). Here, two lines are identical in length. But this is not how they appear. Typically, if you ask someone (or, perhaps, someone who has been raised in a ‘carpentered environment’: Segall et al. 1963), they will report that the bottom line (the one with arrows pointing inwards) looks longer than the top line. While there is disagreement as to precisely why this should be (contrast: Day 1988; Gregory 1997; Sekuler and Erlebacher 1971) things continue to appear this way even when the perceiver knows that this is not so. So, even when the perceiver gets out their tape measure and confirms that the lines are identical in length (and even reflects on this fact) the illusory percept persists.


Fodor’s conjecture, that perceptual processing is composed of informationally encapsulated modules, makes easy sense of this.⁴ On this view, the un-interpreted light hitting the retina underdetermines the length of the lines; it is consistent with the lines having an equal length (after all, they do), but it is also compatible with the bottom line being longer (and, say, further away) than the top line (Gregory 1997). As such, the relevant visual module, tasked with estimating line length, must estimate which of these possible interpretations is most likely. But assuming that it is informationally encapsulated the module can only draw on information stored within its proprietary database to inform its estimation. And, on the conjecture that it is encapsulated from central cognition, this proprietary database fails to include the beliefs and expectations of the perceiver, including those concerning the lines’ lengths. As such, it is unsurprising that knowing the lines to be of equal length has no (apparent) bearing on the way they are perceived to be.

2.3 Encapsulation despite cognitive penetration?

Having introduced my terms, it may seem inevitable that cognitive penetration would be problematic for proponents of Fodorian modularity. For them, perceptual processing is composed of nothing but encapsulated modules, where these modules interpret their inputs on the basis of prescribed proprietary databases that fail to include cognitive information (e.g. the organism’s beliefs, intentions and expectations). In this respect, they are impervious to cognition. But, if cognitive penetration occurs, then cognitive states exert direct and semantically coherent effects on perceptual processing. Many proceed as if this would refute the hypothesis that perceptual processing is composed of nothing but encapsulated modules (e.g. Wu 2017); that, at best, we would be left with “scattered islands of modularity” (Prinz 2006, p. 22).

³ Strictly, this allows for degrees of encapsulation. I will keep matters simple by assuming that perceptual systems are only ‘encapsulated’ from cognition if their proprietary databases fail to include any cognitive information whatsoever. This is often assumed. Allowing perceptual systems *some* access to cognition would only strengthen my case.

⁴ This example is widely discussed, and Fodor’s interpretation is controversial (Prinz 2006). I am simply using the example to illustrate the Fodorian hypothesis.

Fig. 1 The Muller-Lyer illusion

This is a mistake. Admittedly, there is an obvious way in which cognitive penetration might occur that would refute the view that perception is composed of nothing but encapsulated modules: If such systems were to access cognitive states and were to draw on these when interpreting their inputs, they would be systematically affected by cognition in direct and (let us suppose) semantically coherent ways. Thus construed, cognitive penetration would, indeed, call into question the suggestion that perceptual processing is composed of nothing but encapsulated modules, since it would amount to finding that cognitive states enter into the bodies of information that inform (certain) perceptual processes in their inferential analyses. This is just one possibility, however. And, as it so happens, there are other ways in which cognitive penetration could occur. At least one of these would be completely irrelevant to the view that perceptual processing is entirely composed of encapsulated modules, no matter how pervasive it might be.

To illustrate, recall that perceptual processing could comprise nothing but encapsulated modules without each sensory modality constituting a single encapsulated module (something which was never actually Fodor's "intended doctrine"). It could involve each sensory modality comprising a hierarchy of encapsulated modules. This is intelligible because (as we have seen) each module's encapsulation pertains to a restriction on the proprietary database that informs it in the interpretation of its inputs. So, provided that modules within a hierarchy could be identified as perceptual [by appeal to, say, function (see Marr 1982)] and each mechanism within each hierarchy continued to process its inputs on the basis of its own proprietary database (and this alone), the sensory modality would comprise nothing but encapsulated modules, even if the outputs of modules at one level were taken as inputs by modules at the next level of processing.

On a picture of this sort, we could allow that the outputs of a module, located at one level in the hierarchy, be modified by some cognitive process, or even supplemented by the cognitive states of the organism, before serving as inputs to the modules located at higher levels in the hierarchy. Provided that this simply occurred at the joints between modules, and did not add to the bodies of information guiding these systems in the interpretation of their inputs, it would leave the encapsulation of each module intact (i.e. it would not bear on the claim that each module's inputs be processed on the basis of a prescribed body of information—a *proprietary database*—that fails to include cognitive information outside of the system). But since these effects would occur within a functionally defined perceptual hierarchy (and would, presumably, lead to downstream effects of a semantically coherent nature) it would be accurate to speak of perception being cognitively penetrated.

Thus, cognitive penetration of this sort would be perfectly consistent with perception comprising nothing but encapsulated modules, no matter how pervasive it might be.

One might object that a cognitive state's supplementing or modifying the inputs to a module would, strictly speaking, amount to the cognitive state becoming accessible to a module, and (thus) part of its proprietary database. As such, one might worry that cognitive penetration of the aforementioned sort would still violate the encapsulation of (at least some) perceptual modules. This would be a mistake. Inputs and proprietary databases are functionally distinct. The inputs to a module are the ever-changing states (prototypically derived from sensory organs) that the module functions to analyse, while its proprietary database is (by hypothesis) the fixed body of information which informs its analyses. Since the aforementioned type of cognitive penetration simply involves a modification or supplementation of the already changeable inputs to our perceptual modules, it is irrelevant to the claim that modules are restricted to analysing such inputs on the basis of prescribed bodies of information (proprietary databases) which fail to include the cognitive states of the organism.

One might respond that cognitive penetration of this sort still implies unencapsulated computational steps in the perceptual process. This might seem inevitable because the input–output function of the modification/supplementation computation (occurring at the joints between modules) is still mediated by access to cognitive information located outside of perception. But notice: this would only be true insofar as it would be apt to view these modifications/supplementations as *perceptual* computations. And, in at least certain cases, this would plainly be inappropriate. For example: much has been made of the apparent fact that the intentional formation of a mental image involves cognitive influences on our perceptual machinery (see §3). However, a key reason why this has intrigued philosophers and cognitive scientists is the manifest fact that intentionally forming and manipulating mental images seems not to be a *perceptual* process. Rather, it appears to be a kind of imagistic *thought* or *cognition*. So, provided that the modification or supplementation of modular outputs was like this (i.e. a *non-perceptual, cognitive* process) and such influences failed to impact our perceptual processes except by modifying/supplementing the representations being passed between encapsulated modules in a hierarchy of the above sort, effects of this sort would fail to put pressure on the suggestion that every perceptual computation is carried out by an informationally encapsulated module. In this respect, perceptual processing could remain an entirely modular affair, involving nothing more than computations performed by informationally encapsulated modules, despite pervasive cognitive penetration.

Those sympathetic to Fodorian modularity have done much to obscure this point. Leading advocates of this and related suggestions (Firestone and Scholl 2016, p. 4; cf. Fodor 1983, p. 47) have often spoken as if there is a single module corresponding to each Aristotelian sense (thus: a single visual module with a dedicated proprietary database, a single auditory module with a dedicated proprietary database, and so forth) plus or minus one for speech perception. This invites an equivocation between the claim that a sensory modality, like vision, is cognitively impenetrable

and the claim that a given module or computation sub-serving this modality is cognitively impenetrable. But, given the above, we should recognise that this is a mistake: it is only the latter claim that the encapsulation of each perceptual computation implies or requires.

Fodor should have acknowledged this. His exemplar of a modular process was speech perception (1983, p. 64). For him, the modules of speech perception were encapsulated from both central cognition *and* the sensory modules realising vision, audition, and so forth. But this did not involve them receiving inputs from private sensory transducers. Rather, the suggestion was that the initial stages of speech perception take as inputs the outputs of prescribed sensory modules—e.g. visual and auditory modules (McGurk and McDonald 1976)—and process these on the basis of their own proprietary databases (see Liberman et al. 1967, who Fodor cites with approval; see also Liberman and Mattingly 1985, who cite Fodor with approval). So, while it is often suggested that sense modalities, like vision, might each constitute individual modules, we should accept as conceivable the suggestion that perception (within a given modality) be subserved by a hierarchy of fully-encapsulated modules, allowing that a subset of these provide inputs to other modules in the hierarchy (and, ultimately, central cognition). On this view, these inputs could be modified by the organism's cognitive processing, or even supplemented by the organism's cognitive states, at the points where representations pass between modules in the hierarchy, resulting in direct and semantically coherent effects of cognition on perception. But this would not compromise the suggestion that each of these modules be completely encapsulated from all others, and it would not compromise the suggestion that every perceptual computation be carried out by an encapsulated module. So, in short: cognitive penetration of this sort would be perfectly consistent with the conjecture that perceptual processing is an entirely modular affair.

2.4 Does this change the subject?

Cognitive penetration of the above sort is consistent with the informational encapsulation of each and every perceptual mechanism and each and every perceptual computation. It is, thus, consistent with perceptual processing remaining an entirely modular affair. In spite of this, some might object that, in reconciling these phenomena, I have left the modularity hypothesis without the significance it was originally meant to have for cognitive science and the philosophy of perception. Or worse, that I have simply changed the subject.

Such a charge is, I suspect, most likely felt by those who—in objecting to the modularity of perception—were really objecting to the idea that perception as a whole, or perhaps individual perceptual modalities, constitute individual modules. But recall, this was never Fodor's "intended doctrine" (1983, p. 47). And, when we examine the reasons why Fodor and other theorists were originally interested in perceptual modularity, it is non-obvious why the aforementioned proposal should undermine these.

To illustrate: an initial reason why perceptual modularity was originally expected to matter was for its ability to mark a perception-cognition border, thereby

vindicating a folk distinction between perception and thought (Fodor 1983, p. 70; p. 101; 1985, p. 3). For Fodor, this stemmed from the suggestion that central cognition is ‘isotropic’ and ‘Quinean’; *isotropic* in that anything one knows or believes can be brought to bear in central cognitive processing (1983, p. 105) and *Quinean* in that central cognitive inferences are sensitive to the organism’s entire web of belief (pp. 107–8). Both claims are, of course, controversial (see Lewis 1982; Mandelbaum 2016; Norby 2014). But, putting them together, Fodor held that (in principle) anything one knows or believes could be recognised as relevant to their reasoning on a given matter and brought to bear as a premise therein. Since this was tantamount to saying that such reasoning does not operate on the basis of a prescribed and architecturally determined proprietary database it amounted to a rejection of central cognition’s informational encapsulation. But, since cognitive penetration of the above sort leaves the claim that perceptual computations proceed on the basis of prescribed proprietary databases unscathed, it allows that perceptual modules remain encapsulated in a way that central cognition (as Fodor characterised it) is not. Thus, the proposal sketched in §2.3 does nothing to undermine the thought that informational encapsulation marks perception out from central cognition, as Fodor understood matters.

A second reason why the informational encapsulation of perceptual systems was originally expected to matter was for enabling perceptual processes to remain computationally tractable (Fodor 1983, p. 128). Since encapsulated modules interpret their inputs on the basis of dedicated proprietary databases, they need only consider a constrained amount of information when interpreting their inputs. As such, it is relatively easy to see how these systems might sift through all of the information available to them, or organise it in a timely manner, so as to appreciate how it bears on the interpretation of their inputs. But since this (*again*) derives from these systems’ operating on the basis of constrained proprietary databases—that which my proposal leaves intact—informational encapsulation would (*again*) seem to retain the significance it was originally expected to have.

A third reason why the modularity of perception was originally expected to matter was for its epistemic implications (Fodor 1984). From my perspective, this is a more difficult case. Cognitive penetration of the above sort does raise the worry that perceptual justification might (sometimes) be circular (Siegel 2012) and/or liable to influence the reliability of our perceptual faculties (Lyons 2011). Nevertheless, it would be wrong to dismiss my proposal on these grounds.

For Fodor (1984), informational encapsulation bore epistemic significance because it enabled perceptual modules to provide a ‘theory-neutral’ foundation for empirical knowledge. This was meant to derive from the fact that informationally encapsulated modules interpret their inputs on the basis of proprietary databases, free of the organism’s beliefs and expectations. But this always allowed that the organism might wilfully affect the inputs these systems receive (after all, Fodor always acknowledged that a suitably motivated subject might freely avert their eyes or intentionally shift their attention). So, while cognitive penetration (of the above sort) may introduce certain kinds of epistemic worry, it is unclear that an informationally encapsulated system’s operations would not remain ‘theory-neutral’ in much the way Fodor claimed. After all, it would leave each module’s proprietary

database free of background beliefs and expectations. And, while this may fall short of enabling the foundation for empirical knowledge that Fodor had intended, it could retain significant epistemic import (for instance, by constraining imagistic thought in helpful ways).

Of course, Fodor's claims about theory neutrality were always in dispute. Churchland (1988) was quick to object that an informationally encapsulated module would still process its inputs on the basis of a theory, and one which could be just as distorting of the facts as any other; its encapsulation would simply ensure that this theorising remain immune to certain kinds of rational revision. But note: if this is correct, it simply shows that informational encapsulation never really had the epistemic implications Fodor intended. And, once again, this would be true, irrespective of the proposal I am making.

Given disputes of this nature, contemporary proponents of perceptual modularity often distance their hypothesis that perception is modular from specific claims about perceptual epistemology (e.g. Quilty-Dunn 2020). This is not to suggest that there are *no* epistemic implications of their hypothesis. It simply underscores the fact that any epistemic upshots are just this; they are *upshots* of an independent architectural proposal. Thus, we can debate the informational encapsulation of perceptual systems and, in addition, we can debate the implications this would have for the epistemology of perception. The important point (for our purposes) is that these debates are distinct. Consequently: when we examine the implications that perceptual modularity was originally expected to have when Fodor introduced his hypothesis we find little reason for thinking that the proposal sketched in §2.3 robs this of its significance, or that it simply constitutes a change of subject. Instead, we should recognise that the informational encapsulation of each and every perceptual sub-system, and each and every perceptual computation, is consistent with at least one type of cognitive penetration.

3 The empirical point

The preceding discussion makes a conceptual point—it proposes that, in principle, one type of cognitive penetration would be unproblematic for the encapsulation of each and every perceptual sub-system and for the claim that each and every perceptual computation is carried out by an informationally encapsulated module. But this simply charts logical space—it does not tell us anything about how human minds are actually structured. Thus, one might concede the above as a conceptual possibility but doubt that real-life cases of cognitive penetration (assuming they exist) will prove consistent with the informational encapsulation of perceptual systems/sub-systems, let alone the claim that every perceptual computation is carried out by an informationally encapsulated module.

Turning to my empirical point, I now suggest otherwise. Once we recognise the above, I propose that particularly plausible cases of cognitive penetration are both consistent with the informational encapsulation of all perceptual computations and may actually evince this.

Of course, an enormous number of empirical findings have been seen (by some) to demonstrate perception's cognitive penetrability.⁵ It is beyond the scope of this paper to consider each of these in turn. Thankfully, this will not be necessary to make my point. In recent debates, mental imagery has been seen to provide a particularly plausible illustration of perception's cognitive penetrability. For instance, in an influential critique of the modularity hypothesis, Jesse Prinz (2006) claims that mental imagery is “the most obvious” example of cognitive penetration, while Ned Block (2016) has described it as “the most dramatic of the known top-down effects of cognition on perception” (p. 21; see also: Block 2021). This is something both theorists take to undermine or refute the informational encapsulation of perceptual systems (see also: Dijkstra et al. 2019; Howe and Carter 2016; Kosslyn et al. 2006). So: while we have already considered the *possibility* that cases of this sort might prove consistent with the full-blown modularity of perception, showing that these (alleged) cases of cognitive penetration are *actually* consistent with, and even *indicative of*, the proposal sketched in §2.3 would be a significant win for proponents of perceptual modularity. This is what I will now endeavour to show.

3.1 Mental imagery

Several lines of evidence have been advanced in support of the suggestion that mental imagery involves or constitutes cognitive penetration. Some appeal to neural findings suggesting that perception and cognitively formed mental images recruit overlapping brain regions (Block 2021; Dijkstra et al. 2019; Howe and Carter 2016; Kosslyn et al. 2001; Ogilvie and Carruthers 2015). *Prima facie*, this may lend credence to the suggestion that mental imagery can (or does) involve the cognitive hijacking of perceptual resources. But, alone, it is unconvincing. For one thing, proponents of cognitive impenetrability have long been aware of these overlapping neural regions and deemed them irrelevant on the grounds that one cannot draw straightforward conclusions about the cognitive level from findings pitched at the level of neural implementation (e.g. Pylyshyn 2002).

Additional evidence comes from behavioural studies, however. Here, proponents of cognitive penetration find evidence that cognitively driven mental images can become integrated into perceptual processing (Block 2016; Kosslyn et al. 2006). For instance, in an experiment emphasised by Block (2016), Brockmole et al. (2002) gave subjects a “locate the missing dot task,” requiring them to identify a single missing dot from a 5×5 array. Previous work had shown that if a segment from these arrays (containing 12 dots) appeared briefly, within 50 ms of a second segment from the same array (containing 12 separate dots), visual processing would combine the percepts of these segments, enabling subjects to accurately locate the missing (25th) dot. However: if the second segment appeared 100 ms after the first subjects would perform poorly—since this would now exceed the timescale of

⁵ Yale's *Perception and Cognition Lab* lists 182 independent studies, published in 42 different journals, since 1995, purporting to do just this (see <http://perception.yale.edu/Brian/refGuides/TopDown.html>).

visible persistence, they would rely on working memory to represent the first segment and would (consequently) only keep track of (roughly) four dots from it (see: Loftus and Irwin 1998). Amazingly, Brockmole et al. showed that, in spite of this, when the second segment appeared 1500–5000 ms after the first, and subjects were told to try and “imagine the dots still being present after they had disappeared” (p. 317), they would reliably identify the missing dot (almost as reliably as if all 24 dots were presented at once). This is significant for while 1500 ms is considerably longer than the 50 ms the first segment would have visibly persisted for, it is independently predicted to be the length of time it takes to construct a mental image (Kosslyn et al. 2006). So, according to the authors of the study, this provides reason to think a mental image was constructed in light of subjects’ beliefs about the task and was then combined with their percept of the second segment. As such, Block claims that these “results demonstrate a direct content-specific effect of cognition on perception” (2016, p. 22). That is: a case of cognitive penetration as characterised in §2.1.⁶

Of course, there remains room for doubt. For one thing, the mere fact that a percept can be combined with a mental image fails to show that these are combined *within* perception—after all, we think about the things we see all the time. Nevertheless, the suggestion can be bolstered. For instance, in a forthcoming book, Block appeals to a study by Pearson et al. (2008), examining the effects of mental imagery on binocular rivalry (see also Ogilvie and Carruthers 2015).⁷ Here, subjects had conflicting visual patterns presented to each eye. Under normal conditions this leads to a competitive interaction. That is: one pattern becomes dominant in vision for a given period of time, after which the dominant pattern (irresistibly) flips. Pearson et al. showed that merely imagining one of the two patterns facilitated dominance during subsequent presentations (p. 982). This revealed that mental imagery had a similar effect to cases of ‘weak perception’, facilitating dominance in the way an imagined pattern would have had it been presented to the subject at low (but not high) luminance levels (p. 983). This is significant since the functional equivalence found here provides evidence that cognitively driven mental images interact with the same visual mechanisms as those involved in the determination of visual dominance and influence their operations in the same way as weak percepts derived from the retina. So, when considered

⁶ It is worth noting that while this would qualify as a case of cognitive penetration by the criteria laid out in §2.1, stronger characterisations of cognitive penetration may preclude this. For instance: while the (alleged) influence of cognition on perception is ‘direct’ in the sense that each stage in the causal process occurs internally to the organism and at a psychological level of description, it is ‘indirect’ in the sense that the effect of the organism’s cognitive attitudes is mediated by an intermediary (though still cognitive) state (a mental image) (see Macpherson 2012). For some, this may prevent it from qualifying as a case of genuine cognitive penetration. But, however appropriate this stronger notion of ‘cognitive penetration’ might be, it is irrelevant to our concerns here. This is because the (alleged) threat to encapsulation stems from fact that the state in question (the mental image) exists outside of the relevant perceptual modules’ proprietary databases yet seems to be influencing perceptual processing in relevant ways. It is thus, irrelevant for present purposes that it only be an intermediary state, caused by the organism’s cognitive attitudes, causing the perceptual modification in question (as opposed to, say, a belief of the organism).

⁷ Block’s book is still in progress and it is possible that his thoughts on these matters will continue to develop. I am, thus, extremely grateful to him for allowing me to discuss these matters here.

alongside the Brockmole et al. study (mentioned above), these findings provide grounds for thinking that mental images enter into genuine perceptual processing, modulating high-level visual functions, thereby supporting the view that mental imagery can (or does) constitute cognitive penetration.⁸

Now, let us suppose this is so, if only for the sake of argument. My concern is how this should bear on the hypothesis that perception is modular. Block (2021), Prinz (2006) and others (Firestone and Scholl 2016; Ogilvie and Carruthers 2015; Pylyshyn 2002; Robbins 2009) speak as if it would problematize—or even refute—the view. But, given the conclusions of §2, this seems to be a mistake. Any cognitive penetration involved here is, in fact, consistent with the informational encapsulation of each and every perceptual sub-system and would even seem to indicate as much.

To illustrate, note that Brockmole et al. take their findings to vindicate the hypothesis that mental images and visual percepts are combined in a ‘visual buffer’ (p. 329). The visual buffer is a hypothesised functional space onto which mid-level percepts and mental images can be projected and maintained (Kosslyn 1980). Once on the buffer these representations are treated equivalently by subsequent processing (e.g. high-level perceptual systems involved in the determination of visual dominance).

Importantly, this is meant to explain how mental images and percepts can become integrated within the perceptual hierarchy, despite the processes involved relying on independent bodies of information and requiring functional independence in their operations (see Kosslyn 1980, ch. 2). But, given §2.3, this renders the above findings consistent with the full-blown encapsulation of each and every perceptual computation. Since the visual buffer is effectively a joint between functionally independent systems in the visual hierarchy (housing the outputs of *independently posited* low/mid-level visual systems, before these get taken as inputs by independently posited high-level systems [see Marr (1982), and, indeed, Prinz (2012)] the integration of percepts and mental images *here* leaves untouched the idea that every perceptual computation produces its output on the basis of a prescribed proprietary database which fails to include cognitive information. So, even if the integration of mental images and percepts on a visual buffer constitutes genuine cognitive penetration, it is not obvious why this threatens the informational encapsulation of each and every perceptual sub-system/computation in the human mind. It is, thus, consistent with perceptual processing remaining an entirely modular affair.

Pearson et al.’s findings support this contention. Recall: *the* reason why they indicate that mental imagery becomes integrated with perception *in* perception is

⁸ The suggestion that Pearson et al.’s (2008) bolsters Block’s (2016) suggestion that Brockmole et al.’s (2002) evinces cognitive penetration, mediated by mental imagery, is not to deny important differences between these studies. Brockmole et al.’s findings involve cognitive penetration insofar as a (cognitively driven) mental image is superimposed on a percept within the perceptual hierarchy. The conjecture that Pearson et al.’s findings evince cognitive penetration does not turn on any such superimposition—it simply turns on the suggestion that the mental image is influencing a perceptual process (the determination of visual dominance) in relevant ways.

that they purport to reveal the functional equivalence of mental imagery and cases of actual perception (specifically, cases of ‘weak perception’ in the determination of visual dominance). Pending a better explanation, this suggests that mental imagery interacts with the same perceptual mechanisms as genuine percepts when perceptual dominance is being determined. But note: when it is a fully formed percept (derived from the retina) that is interacting with these mechanisms, it would be uncontroversial to suppose that the percept (or some representational precursor) is simply serving as an input to the relevant sub-system (after all, it is derived from the retina). So, to the extent that mental images are processed in a functionally equivalent manner (i.e. the datum that supports their influencing perception proper) we should take this to be evidence that, in mental imagery, mental images are simply serving as inputs to the relevant mechanism(s) in much the same way (if they were not, we would expect to find functional differences between the cases and the argument for cognitive penetration would not go through). But this leaves Pearson et al.’s findings consistent with the informational encapsulation of *all* perceptual mechanisms involved, even if we grant that cognitive penetration is genuinely occurring. Since the mental images in question simply form inputs to the relevant sub-systems determining dominance, and do not bear on the scope of their proprietary databases, these findings leave the encapsulation of these sub-systems intact, in line with the proposal sketched in §2.

This seems to be more than a bare possibility. If a functional space, like the visual buffer, is *the place* where mental imagery becomes integrated into perceptual processing [as the aforementioned studies are meant to evince (see also: Kosslyn et al. 2006)], then we should want to know what it is that ensures that mental images are projected onto the visual buffer and not other locations in the perceptual hierarchy. Here, a natural and (by the modularist’s lights) independently motivated suggestion is that it is the relevant systems’ informational encapsulation that ensures this. On this view, the relevant perceptual modules (those producing outputs on the buffer, and those taking the contents of the buffer as input) process their inputs (derived from the retina, visual buffer or wherever) on the basis of architecturally prescribed proprietary databases that prevent the admission of mental images (created on the fly) or cognitive states that lead to their creation. Consequently, there is simply nowhere else for the mental image to slot into, and thereby penetrate, perceptual processing, other than at the joints between modules (e.g. the visual buffer). Pending a better explanation, these considerations suggest that the above findings are not only consistent with the informational encapsulation of every perceptual sub-system but are indicative of this.

Opponents might respond by denying the existence of a visual buffer, as Kosslyn and others conceive of it. Instead, they might propose that mental images (or the cognitive states driving their production) enter into the bodies of information that guide our perceptual computations in the interpretation of their inputs, rather than just supplementing or modifying their inputs. However, it is the hypothesis that there is a visual buffer, onto which mental images are projected, that the aforementioned studies (cited by my opponents) are designed to test and evince. Indeed, the argument for cognitive penetration based on these studies relies on this. After all, in the Pearson et al. study discussed above, it is the functional equivalence

of mental imagery and weak perception that indicates that mental imagery is modulating genuine perceptual processing. But, as we have just seen, this implies that mental imagery merely forms an *input* to the affected sub-system(s), just as the visual buffer hypothesis predicts.

A more promising response may involve finding independently motivated explanations for integration on the visual buffer that do not appeal to the informational encapsulation of the relevant perceptual systems. Quite what these explanations might be remains unclear, however. Opponents of informational encapsulation have provided many reasons why unencapsulated processes might *appear* encapsulated, and uninfluenced by background knowledge of some variety, despite possessing access to this information. However, standard explanations of this sort seem irrelevant here. For instance: in response to the (apparent) judgement independence of the Muller-Lyer illusion (see §2.2), opponents of perceptual encapsulation might claim that this simply reflects the fact that bottom-up sensory processing trumps top-down predictions within the perceptual hierarchy (Prinz 2006), or that it merely reflects the sluggish speed at which cognitive influences occur (Nicholas Shea, pers. comm.). But neither of these proposals explicates the phenomena currently under discussion. In binocular rivalry and perceptual integration, the mental image has (allegedly) already entered into perceptual processing when it elicits an effect on visual processing, and an appeal to encapsulation accommodates (and, as we have seen, elucidates) this. Thus, the speed at which it enters into perceptual processing is irrelevant. Moreover, it is not the case that the cognitive effect is being trumped by bottom-up visual processing since (by hypothesis) the mental image—once formed—is being integrated into bottom-up visual processing and treated ‘equivalently’ to it. So, pending a better (and so-far lacking) explanation, alleged cases of cognitive penetration mediated by mental imagery could actually seem to indicate that perceptual sub-systems are informationally encapsulated modules.

3.2 Critiquing the view

We have considered cases of mental imagery which have been seen to provide particularly plausible examples of cognitive penetration. But any cognitive penetration they involve seems to be both consistent with, and even indicative of, a view on which perception is entirely composed of informationally encapsulated modules, and on which every perceptual computation is encapsulated from central cognition. I say this because the (alleged) penetration simply concerns a *cognitive* (non-perceptual) process (the formation of a mental image) supplementing or modifying representations at non-arbitrary points in the perceptual hierarchy [specifically, on a visual buffer, housing the outputs of independently posited intermediate-level perceptual sub-systems (Kosslyn 1980; Marr 1982)]. As we have seen, this is naturally accommodated, and even explained, by postulating that the relevant perceptual sub-systems (both the low-mid level perceptual systems which produce outputs on the buffer, and the higher-level perceptual systems which take the contents of the buffer as input) are informationally encapsulated modules. Why? Because this would serve to explain why penetration simply occurs on the buffer, at

the intersection between these functionally distinct sub-systems, and not elsewhere in the perceptual hierarchy.

For what it's worth, I am sympathetic to the thought that similar conclusions will apply elsewhere. For instance, various theorists have claimed that certain forms of top-down attention constitute cognitive penetration, influencing processing within the perceptual hierarchy, as opposed to simply influencing inputs to the perceptual hierarchy as a whole (e.g. Block 2016; Lupyan 2015; Mole 2015; Wu 2017). This has been seen to present a further problem for perceptual modularity (cf. Quilty-Dunn 2020).⁹ Yet, since top-down perceptual attention seems unable to select anything more than a tightly constrained set of features, happenings and property-types, which (again) correspond to the outputs of independently posited mid-level perceptual systems (Wolfe and Horowitz 2004), proponents of perceptual modularity may (again) maintain that these effects simply result from *cognitive influences on the joints between* independently motivated perceptual modules—perhaps, the very same joint as penetration mediated by mental imagery, further motivating the thought that this joint is special for lying between fully encapsulated sub-systems.

In both cases, this seems like good news for proponents of perceptual modularity. But since my discussion constitutes a departure from standard ways of talking about encapsulation and its commitments, I wish to note two (non-exhaustive) ways in which critics might respond.

First: opponents might seek evidence that cognitive penetration (also) occurs at arbitrary points in the perceptual hierarchy. By 'arbitrary' I mean points in the perceptual hierarchy that seem unlikely to correspond to outputs of any interestingly domain specific process that might plausibly constitute a functionally specified module. This does not seem to be true of the cases we have considered since the (alleged) cognitive penetration taking place there involved the elicitation of effects at a joint between *independently posited* modules—namely, on a visual buffer, at the joint between mid and high-level visual sub-systems of the kind posited by Marr (1982), Jackendoff (1986), Prinz (2012) and many others. Plainly this is not inevitable, however. To the extent that cognitive penetration could also be shown to elicit effects at arbitrary points in the perceptual hierarchy (i.e. points that are unlikely to correspond to the joints between functionally independent sub-systems) this would undermine my hypothesis.

⁹ Quilty-Dunn's response on behalf of the modularist is congenial to the proposal made in the present treatment. Indeed, while he does not develop the suggestion, Quilty-Dunn (p. 6) briefly acknowledges the possibility that perception might remain modular (through and through) if it were composed of a hierarchy of modules and cognitive penetration simply occurred at the joints between these. His paper develops a quite different suggestion, however.

He defends an account on which top-down attention might directly influence the internal operations of our perceptual modules, without any cognitive states being *accessed* by the relevant perceptual processes, leaving the encapsulation of these systems intact. So, unlike the proposal defended in this paper, Quilty-Dunn's proposal does not take a stand on whether perception (or a perceptual modality, like vision) actually comprises distinct modules. By contrast, the proposal developed in the present treatment does take a stand on this (it posits a hierarchy of encapsulated modules) but does not require that cognition can directly influence the inner workings of our perceptual modules without modifying their inputs. Thus, the proposals developed are doubly distinct.

Second: critics could seek out cases of cognitive penetration which elicit unencapsulated *perceptual computations* (i.e. computations carried out in perception that rely on access to central cognitive states). If successful, this would force a reassessment of the extent to which perceptual computations are informationally encapsulated, and of the extent to which perceptual computations are solely the result of modular operations. But note, in the aforementioned cases involving mental imagery this does not seem to be so. There, mental images may have been created in light of the organism's cognitive states. And this may lead to their *supplementing* or *modifying* low-mid level perceptual outputs. However, this supplementation/modification computation does not seem to be a *perceptual* computation. Rather, it seems to be a non-perceptual *cognitive* process (Johnson-Laird 2001; Kosslyn et al. 2006). In this way, it presents no challenge to the idea that every *perceptual* computation is carried out by an entirely encapsulated modular system and that perceptual processing remains an entirely modular affair.

4 Conclusion

Cognitive penetration is often taken to undermine perceptual modularity. However, this is non-obvious. Cognitive penetration occurs if one's cognitive states exert direct and semantically coherent effects on their perceptual processing (§2.1), while informational encapsulation obtains insofar as a system processes its inputs on the basis of an architecturally prescribed proprietary database that excludes information accessible to other psychological systems of the organism (§2.2). Despite an apparent tension, this allows that cognition might penetrate perception at the joints between sub-systems in a fully modular perceptual hierarchy, despite leaving each and every perceptual module's encapsulation from central cognition intact, and each and every perceptual computation modular (§2.3). Indeed, once this is recognised, particularly plausible cases of cognitive penetration, like those mediated by mental imagery (by some accounts, *the* most plausible examples of cognitive penetration) would actually seem to *evince* the encapsulation of perceptual systems, rather than refute or undermine this conjecture (§3).

References

- Block, N. (2016). Tweaking the concepts of perception and cognition. *Behavioral and Brain Sciences*. <https://doi.org/10.1017/S0140525X15002733>.
- Block, N. (2021). *The border between seeing and thinking*. (book manuscript).
- Brockmole, J., et al. (2002). Temporal integration between visual images and visual percepts. *Journal of Experimental Psychology: Human Perception and Performance*, 28(2), 315–334. <https://doi.org/10.1037//0096-1523.28.2.315>.
- Brown, H., & Friston, K. J. (2012). Free-energy and illusions: The Cornsweet effect. *Frontiers in Psychology*. <https://doi.org/10.3389/fpsyg.2012.00043>.
- Bruner, J., & Goodman, C. (1947). Value and need as organising principles in perception. *The Journal of Abnormal and Social Psychology*, 42(1), 33–44. <https://doi.org/10.1037/h0058484>.

- Bruner, J., & Postman, L. (1949). On the perception of incongruity: A paradigm. *Journal of Personality*, 18(2), 206–223. <https://doi.org/10.1111/j.1467-6494.1949.tb01241.x>.
- Churchland, P. M. (1988). Perceptual plasticity and theoretical neutrality: A reply to Jerry Fodor. *Philosophy of Science*, 55(June), 167–187. <https://doi.org/10.1086/289425>.
- Clark, A. (2015). *Surfing uncertainty: Prediction, action and the embodied mind*. Oxford: Oxford University Press.
- Coltheart, M. (1999). Modularity and cognition. *Trends in Cognitive Sciences*, 3(3), 115–120. [https://doi.org/10.1016/s1364-6613\(99\)01289-9](https://doi.org/10.1016/s1364-6613(99)01289-9).
- Day, R. H. (1988). *Natural and artificial cues, perceptual compromise and the basis of veridical and illusory perception*. Monash University.
- Delk, J., & Fillenbaum, S. (1965). Differences in perceived color as a function of characteristic color. *The American Journal of Psychology*, 78(2), 290–293.
- Dijkstra, N., Bosch, S. E., & van Gerven, M. A. J. (2019). Shared neural mechanisms of visual perception and imagery. *Trends in Cognitive Sciences*, 43(5), 423–434. <https://doi.org/10.1016/j.tics.2019.02.004>.
- Firestone, C., & Scholl, B. J. (2016). Cognition does not affect perception: Evaluating the evidence for “top-down” effects. *Behavioral and Brain Sciences*, 39, 1–72. <https://doi.org/10.1017/S0140525X15000965>.
- Fodor, J. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.
- Fodor, J. (1984). Observation reconsidered. *Philosophy of Science*, 51(March), 23–43. <https://doi.org/10.1086/289162>.
- Fodor, J. (1985). Precis of the modularity of mind. *Behavioral and Brain Sciences*, 8(1), 1–42. <https://doi.org/10.1017/S0140525X0001921X>.
- Fodor, J. (1988). A reply to Churchland’s ‘perceptual plasticity and theoretical neutrality’. *Philosophy of Science*, 55(June), 188–198. <https://doi.org/10.1086/289426>.
- Fodor, J. (2000). *The mind doesn’t work that way: The scope and limits of computational psychology*. Cambridge, MA: MIT Press.
- Fodor, J., & Pylyshyn, Z. (1981). How direct is visual perception? Some reflections on Gibson’s ‘ecological approach’. *Cognition*, 9(2), 139–196. [https://doi.org/10.1016/0010-0277\(81\)90009-3](https://doi.org/10.1016/0010-0277(81)90009-3).
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston: Houghton Mifflin.
- Gregory, R. (1997). *Eye and brain: An essay on faculty psychology*. Cambridge, MA: MIT Press.
- Gross, S. (2017). Cognitive penetration and attention. *Frontiers in Psychology*. <https://doi.org/10.3389/fpsyg.2017.00221>.
- Gross, S., et al. (2013). Problems for the purported cognitive penetration of perceptual color experience and Macpherson’s proposed mechanism. *The Baltic International Yearbook of Cognition, Logic and Communication*, 9, 1–30. <https://doi.org/10.4148/1944-3676.1085>.
- Halle, M., & Stevens, K. (1962). Speech recognition: A model and program for research. In M. Halle (Ed.), *From memory to speech and back: Papers on phonetics and phonology* (pp. 1954–2002). Berlin: Mouton de Gruyter.
- Helmholtz, H. L. F. (1876). On the limits of the optical capacity of the microscope. *Journal of Microscopy*, 16(1), 15–39. <https://doi.org/10.1111/j.1365-2818.1876.tb05606.x>.
- Howe, P., & Carter, O. (2016). Hallucinations and mental imagery demonstrate top-down effects on visual perception. *Behavioral and Brain Sciences*. <https://doi.org/10.1017/S0140525X15002502>.
- Jackendoff, R. (1986). *Consciousness and the computational mind*. Cambridge: MIT Press.
- Johnston, M. (1992). How to speak of the colors. *Philosophical Studies*, 68(3), 221–263. <https://doi.org/10.1007/BF00694847>.
- Johnson-Laird, P. N. (2001). Mental models and deduction. *Trends in Cognitive Sciences*, 5(10), 434–442. [https://doi.org/10.1016/S1364-6613\(00\)01751-4](https://doi.org/10.1016/S1364-6613(00)01751-4).
- Kosslyn, S. (1980). *Image and mind*. Cambridge: Harvard University Press.
- Kosslyn, S., Ganis, G., & Thompson, W. (2001). Neural foundations of imagery. *Nature reviews. Neuroscience*, 2, 635–642. <https://doi.org/10.1038/35090055>.
- Kosslyn, S., Thompson, W. L., & Ganis, G. (2006). *The case for mental imagery*. Oxford: Oxford University Press.
- Lewis, D. (1982). Logic for equivocators. *Nouûs*, 16(3), 431–441.
- Lieberman, A. M., Cooper, F. S., Shankweiler, D. P., & Studdert-Kennedy, M. (1967). Perception of the speech code. *Psychological Review*, 74(6), 431–461. <https://doi.org/10.1037/h0020279>.
- Lieberman, A., & Mattingly, I. (1985). The motor theory of speech perception revised. *Cognition*, 21, 1–36. [https://doi.org/10.1016/0010-0277\(85\)90021-6](https://doi.org/10.1016/0010-0277(85)90021-6).

- Loftus, G. R., & Irwin, D. E. (1998). On the relations among different measures of visible and informational persistence. *Cognitive Psychology*, 32, 451–461. <https://doi.org/10.1006/cogp.1998.0678>.
- Lupyan, G. (2015). Cognitive penetrability of perception in the age of prediction: Predictive systems are penetrable systems. *Review of Philosophy and Psychology*, 6(4), 547–569. <https://doi.org/10.1007/s13164-015-0253-4>.
- Lyons, J. (2011). Circularity, reliability, and the cognitive penetrability of perception. *Philosophical Issues*, 21(1), 289–311. <https://doi.org/10.1111/j.1533-6077.2011.00205.x>.
- Lyons, J. (2015). Unencapsulated modules and perceptual judgement. In J. Zeimbekis & A. Raftopoulos (Eds.), *The cognitive penetrability of perception: New philosophical perspectives* (pp. 103–122). Oxford: Oxford University Press.
- 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. <https://doi.org/10.1111/j.1933-1592.2010.00481.x>.
- Macpherson, F. (2015). Cognitive penetration and predictive coding: A commentary on Lupyan. *Review of Philosophy and Psychology*, 6(4), 571–584. <https://doi.org/10.1007/s13164-015-0254-3>.
- Mandelbaum, E. (2016). Attitude, inference, association: On the propositional structure of implicit bias. *Nous*, 53(1), 629–658.
- Mandelbaum, E. (2018). Seeing and conceptualizing: Modularity and the shallow contents of perception. *Philosophy and Phenomenological Research*, 97(2), 267–283. <https://doi.org/10.1111/phpr.12368>.
- Mandelbaum, E. (2019). Modularist explanations of experience and other illusions. *Consciousness and Cognition*. <https://doi.org/10.1016/j.concog.2019.102828>.
- Marr, D. (1982). *Vision: A computational investigation into human representation and processing of visual information*. Cambridge: MIT Press.
- McGurk, H., & McDonald, J. (1976). Hearing lips and seeing voices. *Nature*, 264, 746–748. <https://doi.org/10.1038/264746a0>.
- Mole, C. (2015). Attention and cognitive penetration. In J. Zeimbekis & A. Raftopoulos (Eds.), *The cognitive penetrability of perception: New philosophical perspectives*. Oxford: OUP.
- Nagel, T. (1974). What is it like to be a bat? *Philosophical Review*, 83, 435–450.
- Norby, A. (2014). Against fragmentation. *Thought: A Journal of Philosophy*, 3(1), 30–38. <https://doi.org/10.1002/tht3.110>.
- Ogilvie, R., & Carruthers, P. (2015). Opening up vision: The case against encapsulation. *Review of Philosophy and Psychology*, 7(4), 721–742. <https://doi.org/10.1007/s13164-015-0294-8>.
- Palmer, S. (1999). *Vision science: Photons to phenomenology*. Cambridge MA: MIT Press.
- Pearson, J., et al. (2008). The functional impact of mental imagery on conscious perception. *Current Biology*, 18(13), 982–986. <https://doi.org/10.1016/j.cub.2008.05.048>.
- Prinz, J. J. (2006). Is the mind really modular? In R. J. Stainton (Ed.), *Contemporary debates in cognitive science* (pp. 22–36). London: Blackwell.
- Prinz, J. J. (2012). *The conscious brain: How attention engenders experience*. Oxford: OUP.
- Pylyshyn, Z. (1999). Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, 22(3), 341–365. <https://doi.org/10.1017/s0140525x99002022>.
- Pylyshyn, Z. (2002). Mental imagery: In search of a theory. *Behavioral and Brain Sciences*, 25(2), 157–182. <https://doi.org/10.1017/s0140525x02000043>.
- Pylyshyn, Z. (2003). *Seeing and visualizing: It's not what you think*. Bradford: MIT Press.
- Quilty-Dunn, J. (2019). Unconscious perception and phenomenal coherence. *Analysis*, 79(3), 461–469. <https://doi.org/10.1093/analys/any022>.
- Quilty-Dunn, J. (2020). Attention and encapsulation. *Mind and Language*, 35(3), 335–349. <https://doi.org/10.1111/mila.12242>.
- Quilty-Dunn, J., & Mandelbaum, E. (forthcoming). Non-inferential transitions: Mental imagery and association. In A. Nes & T. Chan (Eds.), *Inference and consciousness*. Routledge.
- Rescorla, M. (2015). Bayesian perceptual psychology. In Mohan Matthen (Ed.), *The Oxford handbook of the philosophy of perception* (pp. 694–716). Oxford: Oxford University Press.
- Robbins, P. (2009). Modularity of mind. *Stanford Encyclopedia of Philosophy*.
- Segall, M. H., Campbell, D. T., & Herskovits, M. J. (1963). Cultural differences in the perception of geometric illusions. *Science*, 193, 769–771. <https://doi.org/10.1126/science.139.3556.769>.
- Sekuler, R., & Erlebacher, A. (1971). The two illusions of Müller-Lyer: Confusion theory re-examined. *The American Journal of Psychology*. <https://doi.org/10.2307/1421165>.

-
- Siegel, S. (2012). Cognitive penetrability and perceptual justification. *Nouûs*, 46(2), 201–222. <https://doi.org/10.1111/j.1468-0068.2010.00786.x>.
- Smith, A. D. (2001). *The problem of perception*. Cambridge: Harvard University Press.
- Stokes, D. (2012). Perceiving and desiring: A new look at the cognitive penetrability of experience. *Philosophical Studies*, 158(3), 479–492. <https://doi.org/10.1007/s11098-010-9688-8>.
- Stokes, D. (2013). Cognitive penetrability of perception. *Philosophy Compass*, 8(7), 646–663. <https://doi.org/10.1111/phc3.12043>.
- Tye, M. (1995). *Ten problems of consciousness: A representational theory of the phenomenal mind*. Cambridge, MA: MIT Press.
- Wilson, R. A. (2008). The drink you have when you're not having a drink. *Mind and Language*, 23(3), 273–283. <https://doi.org/10.1111/j.1468-0017.2008.00343.x>.
- Wolfe, J., & Horowitz, T. (2004). Which attributes guide the deployment of attention and how do they do it? *Nature Reviews Neuroscience*, 5(6), 495–501. <https://doi.org/10.1038/nrn1411>.
- Wu, W. (2013). Visual spatial constancy and modularity: Does intention penetrate vision? *Philosophical Studies*, 165(2), 647–669. <https://doi.org/10.1007/s11098-012-9971-y>.
- Wu, W. (2017). Shaking up the mind's ground floor: The cognitive penetration of visual attention. *Journal of Philosophy*, 114(1), 5–32. <https://doi.org/10.5840/jphil201711411>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.