Forthcoming in *Mind & Language*

SUBMITTED ARTICLE

Is pain modular?

Laurenz Casser, The University of Texas at Austin
Sam Clarke, University of Pennsylvania & University of Southern California

Correspondence
Laurenz Casser, Department of Philosophy, The University of Texas at Austin, 2210 Speedway, Austin, TX, 78712-1737
Email: laurenzcasser@gmail.com

Funding information

Clarke's contribution to this paper was supported by a MindCORE Research Fellowship at Penn, and a VISTA Research Fellowship at York University.

Abstract: We suggest that pain processing has a modular architecture. We begin by motivating the (widely assumed but seldom defended) conjecture that pain processing comprises inferential mechanisms. We then note that pain exhibits a characteristic form of judgement independence. On the assumption that pain processing is inferential, we argue that its judgement independence is indicative of modular (encapsulated) mechanisms. Indeed, we go further, suggesting that it renders the modularity of pain mechanisms a *default hypothesis* to be embraced pending convincing counterevidence. Finally, we consider what a modular pain architecture might look like, and question alleged counterevidence to our proposal.

KEYWORDS
pain, modularity, encapsulation, placebo analgesia, cognitive penetration

It is the bitterest pain among men, to have insight into much and power over nothing.

Herodotus, *Histories IX, Ch.16*

1. INTRODUCTION

Can we provide a systematic account of the cognitive processing that underpins pain in humans and other organisms? What would such an account look like? While these are foundational questions for a science of pain, they have received relatively little attention and continue to be overshadowed by research into pain's neuroanatomical implementation. Our aim in this exploratory paper is to begin filling this lacuna.

Our discussion centres around two key questions: First, does pain processing *inferentially* interpret its sensory inputs, drawing on endogenously stored content? Contemporary pain scientists assume so, but rarely defend this assumption in detail (e.g., Wiech et al., 2014; Finlay, 2019). We begin to remedy this shortcoming, offering provisional reasons to regard pain processing as inferential in much the way vision is standardly taken to be.

Second, if pain processing *is* inferential, and allows endogenously stored content to inform its sensory analyses, then *what* endogenously stored content is available to be utilised in these inferential operations? In response, we propose that if pain processing is genuinely inferential, it is likely to be underpinned by one or more *informationally-encapsulated pain modules*—systems which lack access to, and are thus encapsulated from, content stored outside dedicated and prescribed "proprietary databases".

In making these suggestions, we extend traditional discussions of sub-personal inference, encapsulation, and modularity, which focus on paradigmatically perceptual systems (Clarke, 2021; Firestone & Scholl, 2016), speech perception (Fodor, 1983; Liberman & Mattingly, 1985), the systems of core/central cognition (Carey, 2009; Cosmides et al., 2010; Scholl & Leslie, 1999), and the systems of motor control (Mylopoulos, 2021). Additionally, we depart from the small body of work which has considered pain's underlying architecture and deemed it cognitively penetrable and non-modular (Gligorov, 2017; Jacobson, 2017; Shevlin & Friesen, 2021).

We proceed as follows: In Section 2, we note that various phenomena indicate an internal logic to the operations of human pain processing and suggest that this is indicative of inferential pain mechanisms at an algorithmic level of analysis. In Section 3, we build on this suggestion, noting that pain exhibits a characteristic form of judgement independence. On the assumption that pain processing is inferential, we argue that this is indicative of a modular architecture. Indeed, we go further: We propose that it renders the modularity of pain processing a *default hypothesis*—a hypothesis to be embraced pending convincing counterevidence. In Section 4, we consider what a modular pain architecture might look like, distinguishing two versions of our hypothesis. Finally, Section 5 considers some alleged counter-evidence to our proposal. We argue that, given the current state of the evidence, none of this alleged counter-evidence undermines our proposal in even its simplest and most committal incarnation.

2. PAIN AS A PROBLEM OF INFERENCE

Pain is typically characterised as a type of "personal experience associated with actual or potential tissue damage" (IASP).[1] Accordingly, emphasis tends to be placed on the *experiential* aspects of pain. However, few object to the suggestion that pain experiences are underpinned by mechanisms and processes which warrant consideration in their own right. So, like other complex experiential phenomena (e.g., conscious visual perception), pains can be investigated by analysing their underlying mechanisms and processes.

Here, there are various questions to ask. For one, we might ask what (if anything) the relevant mechanisms and processes function to do. In other words, we might pursue a

[1] https://www.iasp-pain.org/terminology?navItemNumber=576#Pain

computational-level description of pain processing (Marr, 1982). For present purposes, we follow others in assuming that pain processing as a whole takes sensory information as input and functions to determine if, and in what way(s), these inputs (or their distal causes) are to be deemed painful for the organism (Seymour & Dolan, 2013).[2] This computational-level description is intentionally vague and non-committal, and it could be fleshed out in various ways. What matters for us is that the realisation of this computational-level task implicates *inferential* processes at a lower, *algorithmic level* of analysis (Marr, 1982).

This much is widely assumed by contemporary pain scientists (e.g., Wiech et al., 2014; Wiech, 2016; Büchel et al., 2014; Pagnoni & Porro, 2014; Finlay, 2019). They conceive of pain as a "problem of inference", involving inferential mechanisms akin to those posited on mainstream "constructivist" accounts of visual processing (Seymour & Dolan, 2013, p. 248; Tabor et al. 2017). In so doing, these researchers implicitly reject "ecological" or "embedded" accounts of pain processing, according to which pains result from mechanisms and processes that are not appropriately described as involving inferences over intermediary mental representations (cf. Gibson, 1979; Orlandi, 2014). Unfortunately, the perceived need for inferential pain mechanisms is rarely articulated in detail. But since "inferentialism" is part and parcel of the conjecture we ultimately recommend, let us consider why this position seems reasonable to assume in what follows.

On constructivist accounts, vision faces an underdetermination problem: When an object reflects light onto the retina, the retinal image is consistent with that object possessing infinitely many possible shapes, sizes, locations, and colours. To overcome such underdetermination, the visual system uses endogenously stored information to inform it in a series of non-demonstrative inferences about the shape, size, location, and colour of the distal objects perceived (Gregory, 1970; Rock, 1975; Palmer, 1999).

To illustrate: The visual system has been said to possess endogenously represented information that light (typically) comes from above (Ramachandran, 1988; Rescorla, 2015a).[3] So when the visual system is presented with some sensory input that is consistent with multiple hypotheses—for example, "convex surface and light from above" and "concave surface and light from below"—this stored content can be accessed, and used to inform the system's analysis in content-respecting ways. Thus, it might lead the system to conclude that the former possibility is more likely. In this sense, the visual system would be said to have literally *inferred* the likely cause(s) of its sensory inputs, transitioning between representational states in non-associative, content-respecting ways (albeit imperfectly: it might draw the wrong conclusion, as in cases of illusion resulting from atypical lighting conditions).

Is this true of pain processing? On traditional "perceptualist" accounts of pain, it is natural to think so. On perceptualist accounts, pain serves the biological function of informing organisms about objective states of bodily damage and is strictly analogous to paradigm forms of perception (Armstrong, 1968; Pitcher, 1970). Hence, pain may be regarded as inferential in much the way vision is. For just as a retinal image underdetermines the size and shape of the

---

[2] "Deeming an input painful" need not involve explicit *categorisation* of the input—it may simply involve *treating* an input as a pain. Furthermore, this computational-level description leaves open *what* pain processing seeks to deem painful for the organism (e.g., bodily damage, disturbances that would benefit from protective behaviours, etc.).

[3] See Sun and Perona (1998) and Adams et al. (2004) for further discussion.

objects one perceives, leaving these to be inferred by visual mechanisms, somatosensory inputs to polymodal nociceptors may underdetermine the nature of bodily damage (see Ringkamp et al,. 2013). By parity of reasoning, this suggests that for pain processes (of the sort posited by perceptualists) to reliably identify objective levels of bodily damage, they, too, perform inferences akin to those performed by other perceptual mechanisms.

Of course, this is no foregone conclusion. While inferentialism enjoys widespread support in the study of vision, it has detractors (Hatfield, 2002; Orlandi, 2014). Nevertheless, the suggestion that *pain* is underpinned by inferential mechanisms can be independently motivated. Consider effects of visual feedback on pain experience. Merely seeing one's pained body part can generate an analgesic effect (Longo et al., 2009; Mancini et al., 2011), as can undergoing visual illusions *as of* seeing one's pained body part. This latter point has proven fruitful in the treatment of phantom limb pain by mirror therapy (Wittkopf & Johnson, 2017): amputees suffering phantom pain can enjoy significant relief when a reflection of their undamaged limb creates an illusory percept as of the amputated limb being intact (Ramachandran et al., 1995). Critically, this is a case of "visual referral" (Ramachandran & Altschuler, 2009; cf. Rock & Victor, 1964). Thus, those suffering phantom limb pain are conjectured to enjoy pain relief from mirror therapy because the illusory percept *as of* an undamaged limb (i.e., an objectified perceptual representation) *contradicts* the analyses that pain systems were reaching about the amputated limb, thereby lowering their credence that bodily damage was incurred. Indeed, Armel and Ramachandran (2003) found evidence that this swings both ways: Visual information that bodily damage has been incurred in a limb (i.e., a visual percept as of your hand being stabbed) irresistibly increases the level of bodily damage that is attributed to that limb. What is notable is that these processes have the form of a valid inferential transition: pain systems take on board new representational content (derived from a visual percept) and process it in content-respecting ways when reformulating their hypotheses. Since analogous cases of visual referral, or "cue combination", provide some of the most convincing evidence for inferential processing in other perceptual domains, where inferentialism enjoys widespread support (Ernst & Banks, 2002; Rescorla, 2015b; Trommershäuser et al., 2011; cf. Orlandi, 2014), the perceptualist finds compelling reason to regard pain processing similarly.

Of course, perceptualism about pain is, itself, controversial (Casser, 2020). Hence, it is important to stress that inferentialism does not depend on a commitment to standard perceptualist accounts. On our preferred view, pain processing does not simply aim at ascertaining some state of the body, but at deciding whether bodily states warrant pain onset given the current behavioural context (Seymour & Dolan, 2013). As such, we think pain processing cares not only *what state the body is in*, but also *what should be done about this* (see Martínez, 2011; Klein, 2015; Martínez and Klein, 2016). We will not argue for this suggestion here but wish to note that if anything like it is true, it only strengthens the case for pain's being "a problem of inference" (despite highlighting a disanalogy with paradigm forms of perception).

We say this because the question of whether pain onset is warranted appears *even more* underdetermined by sensory inputs than the question of how much bodily damage has been incurred. For a start, our preferred account does not deny that bodily damage is relevant to pain processing's assessment of the situation. Indeed, we assume it is. As such, our account accommodates perceptual elements in pain processing which function to gather information about objective states of bodily damage. One reason to suspect that these exist concerns the fact that such information-gathering seems to be subserved by dissociable and domain-specific feedback mechanisms (Sandkühler, 2013). The point is just that accounts of our preferred sort posit that pain processing performs *additional* practical inferences, based on (among other things) these inferred levels of bodily damage. Thus, pain systems may ascertain that the organism (e.g., a soldier on the battlefield or a hunted deer) has sustained major bodily damage but conclude that (given the current fight-or-flight situation) pain onset is best delayed to avoid distraction in the face of immediate danger (Beecher, 1959: Wall, 1979).

What is significant is that these additional transitions, again, have the form of a genuine inference. The system takes (perceived/inferred) content about some objective feature of the world (bodily damage) and, further, content about the situation the organism is in, and performs a content-respecting transition to the conclusion that immediate pain onset is or is not warranted. Indeed, if a view of this sort is correct, pain processing would appear doubly inferential: It comprises perceptual sub-processes which make inferences about objective features of the world (e.g. levels of bodily damage), and further mechanisms which perform practical inferences based on (among other things) the conclusions these sub-processes reach.

We do not consider any of this a foregone conclusion. Critics might dispute the suggestion that an appearance of valid inferential processing at the level of observed behaviour provides reason to expect that inferences occur at the level of information processing (as critics have done when critiquing vision's status as inferential; see Orlandi, 2014). Thus, they might dispute the realist attitude we take towards the above accounts, and the inference-like steps they imply (cf., Rescorla, 2015b). Furthermore, critics might question the descriptive adequacy of the accounts under consideration, noting that they have yet to demonstrate the predictive and explanatory power enjoyed by inferentialist accounts of visual information processing. But while these are legitimate and important concerns, which deserve careful consideration, we content ourselves with a modest interim conclusion: Given the evidence adduced, there is plausible reason to expect that pain processing is inferential in much the way visual processing is standardly taken to be (even if the precise computational functions of pain and vision diverge, as our preferred account suggests).

3. MODULARITY AS THE DEFAULT HYPOTHESIS

3.1. Pain's Modularity

If pain processing is inferential, and draws on endogenously stored content when interpreting sensory inputs, a further question arises: What endogenously stored content is pain processing able to make use of in its inferential operations?

This question would not arise if pain processing were non-inferential (Pylyshyn, 1984, p. xvii). On that view, there is no contentful information informing pain processing in its operations; pain processing simply reduces to a complex interplay of environmental, physiological, and functional factors. But, *if* pain is inferential, as we recommend, the above is surely one of the most fundamental questions we can ask about pain's information processing.

In principle, the information that is available to pain processing in its inferential operations could be unbounded. It could include all things known or believed by the organism and its subsystems, including the organism's beliefs about the severity of relevant sensory disturbances and their significance for the organism. It could even include the organism's beliefs about politics, geography, and clownfish (see Fodor, 1983). Indeed, certain Bayesian accounts of pain processing seem to assume that something of this sort is, in fact, the case: that, under appropriate circumstances, *any* background cognitive states could come to inform pain processing in the inferential interpretation of sensory inputs (Ongaro & Kaptchuk, 2019).

Our hypothesis, that pain is underwritten by one or more encapsulated modules, opposes this suggestion. On our account, pain onset is determined by one or more circumscribed systems which take sensory information as input and (collectively or individually) determine if, and in what way(s), this should be deemed painful for the organism, *entirely* based on dedicated and architecturally prescribed *proprietary databases* (Fodor, 1985). Such databases are (synchronically) fixed bodies of information, dedicated to these pain systems' operations and, most likely, theirs alone. Thus, information located "outside" of a pain system's dedicated database (for instance, what is explicitly known or believed by the organism) is strictly inaccessible to the system and incapable of informing it in the interpretation of sensory inputs no matter how useful this information might be for its operations. Pain systems are, thus, *encapsulated* from information outside their proprietary databases, in much the way Fodor (1983, 1985) regarded the mechanisms of vision and speech perception encapsulated from thought.

Consider the *thermal grill illusion* (TGI). The TGI involves the application of interlaced cold and warm bars on a body part. Provided that these interlaced bars are suitably arranged, and their temperature differs by a sufficiently large margin (roughly ≥ 20°C), bars whose independent application is perceived as painless are collectively perceived as producing a painful burning sensation (Thunberg, 1896; Alrutz, 1898). While there is disagreement as to why this should be, candidate accounts suggest that it involves pain processing making a kind of (inferential) error about the nature, source, or significance of the sensory stimulation it receives as input (Fardo et al., 2020). Yet what is known or believed by the organism seems to have no clear effect on the underlying systems' assessment of the situation. This is so even when the knowledge or beliefs would seem directly relevant to pain processing's operations on any plausible account of what the pain system is trying to infer. So, even when we know that the bars used to produce the TGI are completely innocuous, undamaging, and not something we need protection from, pain onset occurs, and pain persists. In this respect, phenomena like the TGI appear to be *judgement independent*.[4]

Analogous forms of judgement independence have long motivated the conjecture that paradigmatically perceptual mechanisms are encapsulated (Fodor, 1983). While controversial

[4] This is not to deny that our beliefs might influence the phenomenology of our overall experience whilst undergoing the illusion. Plausibly, *knowing* the stimulus to be undamaged will affect one's emotional state (e.g., allaying one's fear of the stimulus), and perhaps enable the subject to attend to other things, thereby influencing their overall experience at a given moment. What we are claiming here is just that knowing the TGI to be innocuous has no effect on the presence or absence of the painful burning sensation itself.

As far as we know, the judgement independence of the TGI has not systematically studied. We here rely on personal communication with Patrick Haggard.

(Prinz, 2007), this helps clarify our proposal. On our view, pain is underwritten by one or more informationally encapsulated pain modules. Here, sensory inputs to these modules fail to unambiguously determine whether pain onset is warranted given the function these systems are (collectively or individually) seeking to perform (e.g., ascertaining levels of bodily damage, or ascertaining whether pain onset will aid the organism's survival, etc.). As such, pain modules make inferences about how to interpret their inputs, drawing on whatever information they have access to. But assuming that they are informationally encapsulated, each module can only access information stored within its prescribed proprietary database. And, on the conjecture that such databases are encapsulated from central cognition, these fail to include the beliefs and expectations of the perceiver, including those which concern the nature of the relevant sensory disturbance (or lack thereof). Thus, on the conjecture that pain is modular, it is unsurprising that knowing the TGI's bars to be innocuous has no (apparent) bearing on pain onset.

3.2. The Case for Modularity

Modularity can explicate the judgement independence of TGI onset. If pain processing comprises modular systems which inferentially interpret their inputs *entirely* on the basis of dedicated, prescribed, and proprietary databases, which fail to include the cognitive states of the organism, we would expect that judging (or desiring, expecting, intending, etc.) the stimulus, sensory disturbance, or situation, *not* to warrant pain onset will have no bearing on our pain modules' inferential conclusions. But this point can only carry us so far. There are various architectures which might accommodate the aforementioned datum, and many researchers assume that pain processing is routinely guided by states of central cognition (e.g., Flor & Turk, 2013). Nonetheless, we argue that, upon inspection, the preceding considerations highlight a compelling reason to favour the conjecture that pain processing is modular (pending convincing evidence to the contrary).

The first point to note is that the judgement independence of the TGI is unexceptional. For while it is true that our beliefs and expectations routinely affect the attitudes we take towards our pains, and the emotional experiences which accompany them, it is hard to identify familiar or everyday cases of pain in which our cognitive states straightforwardly affect our pains' sensory qualities *themselves*.

Consider some mundane examples:

> *Papercut*. You feel a pain in your fingertip as you turn the page. It's a papercut—and it stings like Hell. You examine the wound and find that the page has barely penetrated your skin. You think to yourself: "This really shouldn't hurt so much"—but hurt it does.

> *Headache*. You've had a pounding headache all afternoon. Unable to concentrate on your work, you visit your GP. The GP asks various questions and performs a couple of tests before concluding that you have an ordinary tension headache, caused by a tense neck and shoulder muscles. They reassure you that this is very common and has nothing to do with the state of your head. Regardless, your head continues to hurt on the way home.

These cases are unexceptional. They are ordinary cases of pain which many of us experience on a regular basis. What they illustrate is the familiar fact that our thoughts, hopes, and desires tend to have no noticeable influence on pain onset or persistence. For, as these cases illustrate, one can judge (rightly or wrongly) that the extent of one's injury does not correspond to the intensity of one's pain without this leading to an adjustment in pain intensity (Papercut). Similarly, one can know that the felt location of a pain does not correspond to the source of a relevant bodily disturbance without this causing a shift in pain location (Headache). Indeed, it is difficult to think of cases in which merely reflecting on, or forming, a cognitive state leads to a convincing, noticeable, and synchronic influence on a pain's felt character.

Of course, cognitive states do make a difference to the *attitudes* we take towards our pains and our *emotional reactions* to these. Long distance runners, for example, may employ psychological techniques which help them get through periods of excruciating pain and exhaustion that would otherwise seem insurmountable.[5] Conversely, torturers may employ psychological techniques to heighten their victims' sense of threat, further worsening their experience (Linden, 2015). However, these cases seem less well characterised as instances in which cognitive states influence or change the felt sensory character of pains than as instances in which cognitive states influence how people *manage* or *evaluate* their pains, or as instances in which cognitive states influence the emotional reactions they have *towards* their pains. For example, thinking positive or negative thoughts does not seem to literally render our pains more or less intense, though it plainly affects the attitudes we take towards them and/or our emotional states in relation to our these, rendering these pains more or less bearable.[6]

We have not yet drawn any architectural conclusions from this alleged insensitivity to thought. However, it is somewhat surprising. If pain processing functions to accurately *infer* facts about bodily damage and/or facts about whether pain onset will prove beneficial to the organism (see Section 2), then we might expect relevant judgements to sometimes be considered by pain systems and to have *some* noticeable bearing on the outputs of pain processing. But the fact that we can assess the nature and extent of our injuries, and *know* all sorts about them, does not seem to make a difference to the pains we experience. This is true even when the information would seem directly relevant to pain processing's assessment of the situation, and even when this information would seem to rationally demand revision or reassessment of its conclusions. Indeed, the point enjoys generality and applies no matter what one says about the function of pain. The fact is that it is hard to identify *any* cases in which merely forming, or reflecting on, a cognitive state leads to a compelling and synchronic shift in

[5] Thanks to an anonymous reviewer of this journal for emphasising cases of this sort.

[6] Given that emotions, evaluations, and attention are routinely influenced by our cognitive states, we believe that the onus lies with those who reject our introspective assessment of the situation to make the case that things of this sort are *not* all that changes with our shifting judgments during episodes of pain. But even if this challenge is met, our argument for pain's modularity would still apply to certain sub-components of pain processing provided that judgements do not, and cannot, entirely eliminate our pains. For if it were true that judgements can merely influence (e.g.) the felt character of one's pains, but not (e.g.) whether pain onset occurs and pain persists (this latter point seems completely indubitable to us), the following remarks would be taken to suggest that the determination of *pain onset and pain persistence* is subserved by modular mechanisms, even if subsequent unencapsulated mechanisms determine the character of the resulting pain. So, in sum: even if one rejected our phenomenological assessment of the situation in its entirety, we think that one should still find reason to think of pain as *partially modular* in the argument that follows.

a pain's felt qualities. (Readers with a specific view on pain's computational function are invited to consider whatever judgements would seem most relevant on their preferred account.)

Admittedly, and as acknowledged earlier, judgement independence does not *prove* that inferential pain mechanisms are encapsulated. It is, however, a datum that a modular architecture makes straightforward sense of. For if pain processing is underpinned by encapsulated modules, as we conjecture, then we would expect that these pain systems be judgement independent in the aforementioned ways. Meanwhile, those who oppose this suggestion have some explaining to do. Assuming that pain processing is inferential, they must explain why pain systems, which are not encapsulated from cognitive information, routinely fail to let it alter their conclusions, even when this would be useful, appropriate, or beneficial for the fulfilment of their computational function(s) and inferential goal(s). This, we contend, is easier said than done.

One response might be that even if pain systems lack architecturally prescribed proprietary databases of the sort we propose, speed of information flow within the mind/brain prevents cognitive states from reaching pain systems in time to inform them in their inferential interpretations. On this view, pain processing would not be subserved by one or more informationally encapsulated modules, but pain analyses would wind up judgement independent by virtue of our cognitive states/cognitive information taking *too long* to arrive at the relevant systems and, thus, too long to get incorporated into their inferential analyses.

This might seem to provide a neat explanation for the above observations. After all, it does not appeal to the cognitive architecture of pain processing. It merely concerns facts and constraints (speed of information flow) that any account of pain processing must contend with. Yet, on inspection, the suggestion does not provide a plausible explanation for pain's characteristic judgement independence. For while it is possible that the (relatively sluggish) speed of top-down information flow within the mind/brain might offer a non-modular explanation for the judgement independence found in certain cases of pain and certain paradigmatically perceptual phenomena—cases in which sensory systems function to report on what is happening *right now* (Fodor, 1989; Mandelbaum, 2018)—pain onset is often remarkably slow. For example, onset of acute pain is frequently delayed by minutes, if not days (Melzack et al., 1982; Wall, 1979), while pain following a stroke or spinal cord injury can take years to set in (Schott, 2001). Even pain onset in the TGI is typically delayed by several seconds (Fardo et al., 2020). Consequently, one's cognitive assessment of an injury (or lack thereof) often precedes pain onset by a significant margin. So, while it may be important for pain systems to reach prompt and timely interpretations of their inputs in certain situations, it is implausible (as a more general claim) that cognitive states ordinarily have insufficient time to reach pain systems before they complete their analyses. Thus, speed of information flow is not a convincing explanation for pain's characteristic judgement independence.

An alternative response could be that pain's judgement independence results not from the encapsulation of pain systems, but because judgements tend not to alter the resulting experiences for extraneous reasons: For instance, it might be claimed that pain systems only bother to search, access, and consider cognitive states of the organism when they are (otherwise) unsure how to interpret their inputs (Prinz, 2006). Alternatively (or perhaps additionally), it might be claimed that pain processing simply places greater weight on the outputs of low-level pain mechanisms, performing their analyses in a bottom-up fashion (based on sensory inputs), such that these outputs tend to trump conflicting top-down judgements in the overall analysis (Clark, 2013). On either account, it might be claimed that pain systems typically end up displaying the aforementioned kinds of judgement independence, despite their lack of encapsulation. And since accounts of both sorts offer popular explanations for the judgement independence of visual processing among critics of the modularity hypothesis (Howhy, 2013), they may prove attractive to those sceptical of our proposal also. But note two things:

First, in the case of vision, such possibilities are motivated by *actual* examples—e.g. bistable images (like the necker cube and duck-rabbit) which are ambiguous between two (or more) interpretations, and cases of binocular rivalry (e.g., Howhy et al., 2008). In such cases, ambiguity often results in one visual interpretation dominating for some period of time before a gestalt shift occurs and the alternative interpretation comes to dominate conscious experience. In such cases, non-modularists sometimes posit that cognitive states are accessed by visual systems, and able to guide them in their interpretation of the image, disposing one interpretation over the other, owing to the uncertainty of the bottom-up analysis (cf. Block, 2016). Thus, it is sometimes proposed that this is why children are more likely to see the duck-rabbit as a rabbit on Easter weekend than at other times of year (Brugger & Brugger, 1993; Prinz, 2006).

Modularists find such examples woefully unconvincing (Mandelbaum, 2019). But, even if they did offer a compelling violation of *vision's* encapsulation, there is a dearth of evidence for anything like a bistable image in the case of pain: there simply are no cases in which we experience gestalt shifts between (e.g.) *painful* and *less/not-painful,* or *pain in location A* and *pain in location B*. This is not to deny that the inputs to pain processing are, themselves, ambiguous (see Section 2); nor is it to deny that our judgments affect the attitudes that we take towards, or the emotions that accompany, our pains. The point is just that we currently have no reason to think that any of this leads to bistability in the pain itself (let alone bistability that is informed by the organism's cognitive states).

This raises a second and more fundamental issue. The suggestion that cognitive states might be *accessible* yet *not actually accessed* by pain systems, or *accessed* without this influencing the pain system's overarching analysis, amounts to no more than a bare possibility unless we have positive reasons to endorse it. For without reasons of this sort, the informational encapsulation of pain systems should constitute our *default hypothesis*. We say this because it offers to *explain* the aforementioned (and seemingly pervasive) types of judgement independence that are so characteristic of pain; on this view, judgements fail to inform pain analyses *because* they are architecturally inaccessible to pain processing in its inferential analyses. Since the alternative—that these states might be accessible to pain processing, yet remain routinely un-accessed by it, or accessed without this having any effect—merely accommodates the datum, it provides a shallow and seemingly *ad hoc* account of the phenomenon. So, without positive reason to favour some such alternative, pain's modularity should be considered the leading conjecture, given the considerations we have amassed so far.

4. TWO KINDS OF MODULAR ARCHITECTURE

Our exploration of pain's information processing architecture has now taken us from the defence of a widely endorsed claim (that pain is inferential) to the recommendation of a controversial idea (that pain processing's inferential operations are performed by informationally encapsulated modules). But before we consider where critics might push back, it is worth considering what our recommended modular architecture might look like, explicating its commitments more explicitly.

A simple version of our hypothesis would posit one "large" pain module. On this view, there would simply be one proprietary database that component sub-processes involved in pain analyses would be constrained to access and make use of in their computations. Importantly, this would not preclude *certain* outside influences on pain processing. For one, the system's proprietary database would still function to drive the analysis of *sensory inputs* deriving from elsewhere (e.g., sensory transducers and the outputs of perceptual mechanisms, as in cases of visual referral [see Section 2]). Thus, the view would still predict that pain processing be influenced by the outputs of the mechanisms delivering its inputs. But since this is true of modular systems quite generally,[7] and since the inputs to a module and its proprietary database are functionally distinct (*inputs* are the ever-changing states that the module functions to analyse, while the *proprietary database* is a fixed body of information which [by hypothesis] informs these analyses [see Clarke, 2021]), such influences are unproblematic for the view. What would be problematic is if the single monolithic pain module were to access content stored outside its proprietary database, and use this in semantically-coherent ways when analysing its sensory inputs. In other words, it would be problematic for this version of the view if the (singular) pain module was cognitively penetrated in the course of transforming its prescribed inputs into considered outputs (Pylyshyn, 1999).

By contrast, an alternative account of pain's modular architecture would posit *multiple* pain modules. On this view, each module would simply function to perform some dedicated part of the overall pain analysis. Suppose, for instance, that pain systems function to decide whether bodily states warrant pain onset given the current behavioural context (Seymour & Dolan, 2013). As we noted in Section 2, a view of this sort is naturally aligned with the existence of both perceptual elements in the pain process (functioning to infer levels of bodily damage) *and* additional elements, performing practical inferences on the basis of (among other things) the outputs of these perceptual elements. But since it is reasonable to expect that the proprietary bodies of information that would usefully drive relevant perceptual functions (i.e. inferring bodily damage) would be *very* different from those which drive a system's practical assessment of whether pain onset is or is not warranted in a given situation, it might be reasonable to expect that these functions are subserved by distinct modules: a "bodily damage module" and an independent "practical inference module",[8] if you will.

A multi-module architecture of this sort would be far from unprecedented. Marr (1982) posited low, mid, and high-level modules in his influential sketch of visual analysis. And, when Fodor first sketched his modularity hypothesis, he began by rejecting the possibility that there is one big module corresponding to each of the Aristotelian senses, remarking that this was *not* his "intended doctrine" (1983, p. 47). Instead, he proposed that "within (and quite possible

[7] Fodor's (1983, p. 64) exemplar of a modular system was phoneme perception, which (as illustrated illustrated by the McGurk effect) was seen to involve amodal modules functioning to analyse objectified outputs of both visual and auditory analyses (see Liberman et al., 1967, who Fodor cites with approval; see also Liberman and Mattingly, 1985, who cite Fodor with approval)

[8] In talking of a *practical inference module* we mean to denote a module that performs the practical inferences that are proprietary to pain analysis on the aforementioned account. This should be distinguished from the proposal that practical inference, more generally, might be subserved by modular processes (e.g. Carruthers, 2004).

across) the traditional modes there are specialised computational mechanisms", where the specialisation of these mechanisms derives from "restrictions on the information they can access" (ibid.). Thus, we should take seriously the possibility that the same might be true of pain processing.

Consideration of this alternative is important because the commitments of a multi-module architecture are quite different from those of a simpler single-module architecture. Most critically, a multi-module architecture is compatible with the occurrence of cognitive penetration *within* pain processing itself, provided that this penetration simply occurs at the joints *between* independently posited systems, influencing the outputs of lower-level modules before these are taken as input by higher-level systems (Clarke, 2021). This is because (a) we have already seen that the inputs to a module are functionally distinct from its proprietary database, and (b) a violation of encapsulation simply turns on the idea that the system accesses information outside of its proprietary database when analysing its inputs; *ergo* cognitive effects on the system's inputs are irrelevant. So, putting these points together: a multi-module architecture permits cognitive penetration at the joints *between* modules provided that this does not alter the input-output function of independently posited modules within the system.

5. EMPIRICAL OBJECTIONS

We can now take a step back, and better appreciate the commitments of our hypothesis that pain is a modular process. First, we can see that it is consistent with two broad modular architectures which differ in their commitments to the ways in which pain processing might be penetrated by outside information. It is thus a mistake to assume that cognitive penetration straightforwardly implies that pain processing could not be "organised into discrete processing modules" (*pace* Shevlin & Friesen, 2021, p. 773). (For discussion, see Clarke, 2021; Mylopoulos, 2021; Quilty-Dunn, 2020.) Second, we can see that *any* account of pain's modularity allows that both the inputs and outputs of pain processing (as a whole) be modulated by external influences. This is because the encapsulation of a system only pertains to restrictions on the scope of its proprietary database and, hence, information that can be accessed by the system in the course of analysing its inputs and transforming these into appropriate outputs. So, anything that happens before or after this transformation process fails to bear on the proposal in question.

With these points in view, let us now consider some proposed counterevidence to the modularity of pain processing, asking how it bears on our hypothesis. In each case, we argue that the evidence is, in fact, perfectly consistent with the full-blown modularity of pain processing. Indeed, we argue that it is perfectly consistent with the maximally committal, single module architecture outlined above, according to which pain is never penetrated in the course of its inferential operations.

5.1. Attention and Multimodal Perception

Papers emphasising the complexity of pain processing frequently stress the influence of multisensory inputs and top-down attention on pain experience (e.g. Senkowski et al., 2014). In either case, pain analyses are influenced by information located externally to pain processing itself. But while attention (Wu, 2014, cf. Quilty-Dunn, 2020) and multimodal perception

(Nanay, 2018; Prinz, 2006; cf. Burnston & Cohen, 2015) have been deemed problematic for modularity in other domains, such as vision, it is difficult to see why either phenomenon would undermine the encapsulation of inferential pain mechanisms.

On the topic of multisensory inputs: since we have clarified that a system's encapsulation simply concerns the existence of architectural constraints on the information that the system has access to when interpreting their inputs (i.e. on their proprietary database), it is no objection that these inputs might derive from multiple sensory sources (e.g. visual, auditory, and tactile systems, as well as those of pre-perceptual sensory transducers). Indeed, seminal discussions of encapsulation and modularity, like Fodor's (1983), treated systems with access to multisensory inputs, like the systems of phoneme perception, as *the* paradigms of the modular kind (ibid., p. 64). Thus, the influence of multisensory inputs on pain processing is fully consistent with full-blown encapsulation, and it is odd that critics of perceptual modularity frequently assume that this is in tension with the view (e.g., Nanay, 2018; Prinz, 2007).

Attentional control is similarly unproblematic. Evidence does indicate that allocating top-down attention towards the source of one's pain increases its reported intensity (Wiech et al., 2014). Nevertheless, it is obscure to think that this should be interpreted as a case in which pain's inferential processing *accesses*, and is *informed,* by cognitive states of the organism, or information external to the proprietary databases of pain processing. For one thing, the intention to attend to one's pain does not provide inferential pain processing with reason to deem the relevant bodily disturbance *more painful*. So, even if attention systematically influences pain processing in this way, it is obscure to think that it involves a violation of pain systems' encapsulation.[9] Indeed, since it is uncontroversial to suppose that attention influences perception and perceptual processing, it is plausible to suppose that the intention to shift one's attention towards the source of one's pain simply leads to selection and prioritisation of the relevant spatial region in visual/tactile perception, and that it is this process which then modulates the inputs that pain mechanisms receive. So, even if these attentional effects violate the encapsulation of visual/tactile systems (a point which is, itself, controversial), there is no reason why this would (in turn) undermine the encapsulation of pain processing itself.

5.2. Placebo Analgesia

A more targeted source of scepticism about *pain's* modularity concerns placebo analgesia. Critics have recently claimed that this is the clearest and most compelling illustration of pain's cognitive penetrability (Gligorov, 2017; Shevlin & Friesen, 2021). And while their discussions may be less focussed on the precise architectural details of pain systems—as opposed to the epistemological and ethical upshots of cognitive influences on such systems—they have taken cognitive penetrability to be directly relevant to theories of mental architecture (ibid., p. 775), and to indicate that thus affected mental processes cannot be "organized into discrete processing modules" (ibid., p. 773). This makes them an ideal stalking horse for our purposes.

[9] Related points have been made in the perception literature. In an illuminating discussion, Gross (2017) notes that the imperative "look for red things" fails to provide an epistemic reason to *see* red things. In this way, he proposes that the intention to attend to red things cannot literally be informing, and thereby penetrating, visual processes in the way that theorists like Pylyshyn (1999) denied (cf. Green, 2021). Our point is stronger. Not only does the intention to attend to our pain fail to provide any epistemic reason for pain (from the perspective of the pain systems' inferences), it fails to even provide a practical reason for thinking the relevant body part more *damaged*/something that needs *more protection in the here and now*.

Placebo analgesia is apparent pain reduction "that results from a subject's perception of therapeutic intervention, regardless of whether the intervention is an active or inert agent" (Schmidt & Willis, 2007). Thus, it involves cases in which the appearance of an effective treatment (e.g., a tablet, nasal spray, topical cream, acupuncture, or surgery) generates a significant analgesic response (Wager & Fields, 2013). To succeed, conditioning of pain relief with explicit sensory cues, (conscious) expectations of treatment efficacy, and the psychosocial context surrounding treatment are crucial (Price, 2008). Thus, cases of placebo analgesia are taken to suggest that open administration of medical care, and promises of treatment efficacy, can make significant differences to subjects' treatment responses. On a standard interpretation, this is because open treatment creates positive expectations of analgesia, which are linked to pain relief (Wager & Fields, 2013).

Proponents of pain's cognitive penetrability (and non-modularity) take this to reveal that higher-level cognitive states, such as beliefs and judgements directly modulate lower-level sensory representations involved in pain processing (Shevlin & Friesen, 2021). More specifically, these cases of placebo analgesia are taken to indicate an "extreme" or "radical" form of cognitive penetration in which cognitive states exert a synchronic, reason-respecting effect on pain processing, where this does not involve non-psychological steps in the causal process, such as, for example, manipulation of sensory organs or, perhaps, covert shifts of attention (ibid., p. 775). Indeed, those advocating these claims state that the evidence for pain's cognitive penetrability is *so robust* that it renders alleged evidence of visual penetration "scanty" by comparison (ibid., p. 772)!

The difficulty with these assertions is that there are numerous, independently plausible and familiar, ways of explaining the cited cases of placebo analgesia which would not involve any cognitive penetration whatsoever. So, even bracketing the jump from cognitive penetration to non-modularity, which we have questioned, the assumed penetrability of pain processing is dubious at best.

Shevlin and Friesen do acknowledge several ways of explaining away pain's alleged cognitive penetrability in the present context. For instance, they acknowledge that purely associative influences on pain processing can account for the cognitive influences exhibited in some cases of placebo analgesia (Shevlin & Friesen, 2021, p. 778), and note, correctly, that such *purely associative* effects would fall short of genuine cognitive penetration (Pylyshyn, 1984, p. xvii). This is because cognitive penetration requires a quasi-logical, "semantically coherent" effect on pain processing, at least insofar as it bears on the encapsulation hypothesis (Pylyshyn, 1999). Moreover, while Shevlin and Friesen maintain that certain cases of placebo analgesia result from non-associative influences on sensory mechanisms, they recognise that *some* of these may simply reflect decision or response biases ("appraisals")—decisions and responses which would constitute *post-modular* happenings, rather than alterations of the input-output function pain modules compute.

Despite these observations, Shevlin and Friesen insist that certain cases of placebo analgesia *cannot* be explained in any of these ways. They provide two arguments for this claim. Their main argument turns on neural evidence indicating that "expectations related to pain can impact the spinal processing of pain" (2021, p. 780). Since spinal processing is "thought to occur downstream before any cognitive appraisal occurs" (ibid.) and is seen to constitute some of the earliest neural activation involved in pain processing, the authors conclude that this could not simply reflect post-pain-system appraisal or reports of the pain system's outputs.

We find this plausible enough. But note: the appeal to neurological evidence is a double-edged sword from their perspective. To establish that pain processing is cognitively penetrable, we have already seen that it is not enough to show that cognitive states fail to simply influence post-pain-system analyses (i.e. reports or appraisals). It must also be shown that the effects in question are not simply influencing inputs to the entire inferential process. And this should worry our critics. For insofar as they appeal to neural findings, indicative of influences on *really early* sensory mechanisms (e.g. early spinal processing) it is likely that these influences are merely eliciting pre-inferential influences of this sort.

This is a familiar and fundamental point (Deroy, 2013, Firestone & Scholl, 2016); it is why proponents of vision's cognitive penetrability have been at pains to suggest that (e.g.) shifts of spatial attention do not operate in this way (Mole, 2015; Wu, 2014). And, while the inferential architecture of pain processing remains poorly understood, the extant evidence indicates that inferential pain analyses do, in fact, occur at levels of neural organisation much higher than the spinal cord. For instance, Ferrè and colleagues (2018) found evidence that the TGI is resistant to tactile modulation at the level of the spinal cord, indicating that the painful sensation results from inferential processing occurring later, at the level of the cortex.

Shevlin and Friesen's second argument is made in passing. It is based on experimental results from Wiech et al. (2014). Wiech and colleagues used a drift diffusion model to analyse how cognitive states (expectations) might influence pain processing in placebo analgesia. As Shevlin and Friesen rightly note, the authors of this study reported finding that pain processing is sometimes impacted by expectations, which influences "the inferential process underlying [pain] perception in which prior information is used to interpret sensory information" (ibid., p. R679). However, the suggestion that this threatens the encapsulation of pain processing seems to reflect a terminological confusion and a misunderstanding of the model employed.

Drift diffusion models enable researchers to model various factors that contribute to subjects' decision making under tightly constrained conditions (namely, in 2-alternative-forced-choice tasks, where decisions are made promptly, in <1500ms [see Ratcliffe & McKoon, 2008]). By looking at error rates in tasks which require subjects to report whether a painful stimulus was of a high or low intensity, whilst presenting them with visual cues which signalled the presence/absence of a high-intensity pain (with varying reliability), Wiech and colleagues probed the question of whether this probabilistic information affected their post-sensory analyses of the painful stimuli. They found that it did. In fact, they found that the effect of prior information on pain reports was "predominantly based on altered perceptual decision making" rather than altered sensory analyses (R679). Crucial, however, is that the term "perceptual" is something of a misnomer in the present context. For in talking of "altered perceptual decision making" and expectancy-related influences on "the inferential process underlying [pain] perception", Wiech and colleagues simply take themselves to have provided evidence of expectancy-related effects on "perception" in a very loose sense; strictly speaking, they are talking about expectancy-related effects on the judgments subjects form about the percept. Indeed, the authors are explicit on this point, writing that their results challenge the view that pain "perception" is modulated entirely by sensory-discriminative brain regions "as opposed to report bias" (R680). Thus, if we accept their results, Wiech and colleagues have provided evidence that expectancy related modulation of pain "predominantly" reflects post-perceptual decision making (or, in Shevlin and Friesen's terms, post-modular "appraisal") which, as Shevlin and Friesen (rightly) acknowledge, is congenial to the modularist who denies pain's cognitive impenetrability entirely.

Hence, neither of Shevlin and Friesen's arguments poses a threat to our proposal: in fact, the studies adduced vindicate bold predictions of a crude and simple version of the modularity hypothesis. Namely, that the apparent violations of encapsulation (under consideration) simply reflect modulations of the pre-inferential inputs to pain processing, or subjects' post-modular assessments of the situation.

5.3 Practical Assessments

Perhaps a deeper problem for pain's modularity is of our own making. Recall that on our preferred account, pain processing takes into account information about the situation that the organism is in and performs *practical inferences* about whether pain onset would be beneficial to the organism (hence why a soldier on the battlefield may not experience immediate pain onset following a severe injury [Section 2.1]). For the purposes of this paper, this is a conjecture we remain neutral on. Nevertheless, critics might think it presents an additional worry for our hypothesis. For how could the pain system know that the soldier is in a perilous situation, and not simply relaxing at home when their leg is blown off? Does this not require that pain systems have access to their organism's *beliefs* about the situation?

We think not. For one thing, there are many ways that a modular pain system might estimate the severity of the situation using prescribed sensory cues, without drawing on any extra-sensory, cognitive information whatsoever. For instance, heuristic processes might take into account information about the organism's heart rate and stress levels and use this information to track the severity of the situation. For if these sensory inputs correlate with danger (to some degree), their consideration could enable pain systems to track such matters reliably (albeit not infallibly). Indeed, systems of this sort might even use perceived information about levels of bodily damage to predict danger if particularly high levels of bodily damage correlate with dangerous situations. What's critical is that pain systems would not be accessing the organisms' beliefs or cognitive states in either case. And, indeed, there is positive reason to suppose they are not. For while a soldier on the battlefield may know that they are in a dangerous situation (one in which they would benefit from a delay of pain onset), pain onset can also be delayed or "gated" in cases where there is no danger of this sort. For instance, those suffering injuries in the home or workplace frequently find pain onset delayed in comparable ways, even when they know themselves to be in no danger and in a situation where they might benefit from immediate pain onset (e.g. to encourage nursing behaviour—Melzack et al., 1982). So, since the subject's knowledge can conflict with the relevant systems' analysis of the situation, these analyses would, themselves, appear to be judgement independent. Thus, we find reason to posit that tracking the severity of the situation results from judgement independent processes, consistent with (and arguably indicative of) encapsulated mechanisms.

6. CONCLUSION

The cognitive processing that underpins pain in humans (and other organisms) is a neglected topic. The present treatment has sought some progress on this issue by articulating and defending several bold conjectures. Section 2 motivated the (oft-assumed, but seldom-defended) view that pain processing is inferential and draws on endogenously stored information to inform its sensory analyses. Section 3 proceeded to motivate a conditional claim: that, if pain processing is inferential, it likely possesses a modular architecture, comprising encapsulated mechanisms which perform inferential interpretations of their inputs (entirely) based on information in dedicated and prescribed proprietary databases. Indeed, we went further, proposing that this should be considered our *default hypothesis*—a hypothesis to be embraced, pending convincing counterevidence. Section 4 clarified what this might amount to, distinguishing single from multi-module pain architectures. And Section 5 considered a range of alleged and possible counter-evidence to our proposal, finding that it is entirely unproblematic for either of the modular pain architectures described in Section 4.

Despite these modest victories, we do not take any of the above to be a foregone conclusion. Indeed, Section 2 discussed ways in which the case for inferentialism about pain might be questioned, while Section 3 simply argued that relative to *one* source of evidence (judgement independence) modularity offers the most elegant explanation. Even in Section 5, we merely showed that rejections of encapsulation are undermotivated. But, however provisional our conclusions, we are hopeful that our discussion provides a promising framework within which to think about the architecture of pain processing. Rather than proceeding on the assumption that pain processing operates in a radically top-down manner, as many theorists assume from the outset, extant evidence provides reason to proceed cautiously and on the (provisional) assumption that pain is probably a modular process.

ACKNOWLEDGEMENTS

The authors are extremely grateful to audiences at the University of Oxford and the University of Texas at Austin for helpful questions and comments on earlier versions of this paper. We would also like to thank Jake Beck, Will Davies, Gary Hatfield, Dylan Ludwig, Jake Quilty-Dunn, Melissa Rees, and Mark Sainsbury for detailed commentary and discussion.

REFERENCES

1. Adams, W. J., Graf, E. W., & Ernst, M. O. (2004). Experience can change the "light-from-above" prior. *Nature Neuroscience*, *7*(10), 1057–1058.
2. Alrutz, S. (1898). On the temperature-senses. *Mind*, *7*(25), 141–144.
3. Armel, K. C., & Ramachandran, V. S. (2003). Projecting sensations to external objects: evidence from skin conductance response. *Proceedings. Biological sciences*, *270*(1523), 1499–1506.
4. Armstrong, D. M. (1968). A materialist theory of the mind. New York, NY: Humanities Press.
5. Beecher, H.K. (1959). Measurement of Subjective Responses. New York: Oxford University Press.

6. Block, N. (2016). Tweaking the concepts of perception and cognition, *Behavioral and Brain Sciences,* 39, e232.
7. Brugger, P., Brugger, S. (1993). The Easter bunny in October: Is it disguised as a duck? *Perceptual and Motor Skills*, 76(2), 577–578.
8. Büchel, C., Geuter, S., Sprenger, C. & Eippert, F. (2014). Placebo analgesia: a predictive coding perspective, *Neuron*, 81(6), 1223-1239.
9. Burnston, D.C., & Cohen, J. (2015). Perceptual Integration, Modularity, and Cognitive Penetration. In Zeimbekis, J. & Raftopoulos, A. (Eds.), *The Cognitive Penetrability of Perception: New Philosophyical Perspectives* (pp. 123-143). Oxford: Oxford University Press.
10. Carey, S. (2009). The origin of concepts. New York: Oxford University Press.
11. Carruthers, P. (2004). Practical Reasoning in a Modular Mind. *Mind & Language*, 19(3), 259-278.
12. Casser, L.C. (2020). The Function of Pain, *Australasian Journal of Philosophy*, 99(2), 364-378.
13. Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3), 181-204.
14. Clarke, S. (2021). Cognitive penetration and informational encapsulation: Have we been failing the module?, *Philosophical Studies*, 178(8), 2599-2620.
15. Cosmides, L., Barrett, H. C. & Tooby, J. (2010). Adaptive specializations, social exchange, and the evolution of human intelligence. *Proceedings of the National Academy of Sciences*, 107(Supplement 2), 9007–14.
16. Deroy, O. (2013). Modularity of perception. In Matthen, M., (Ed.), *The Oxford handbook of philosophy of perception* (pp.755–780). Oxford: Oxford University Press.
17. Ernst, M.O. & Banks, M.S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature,* 415, 429-433.
18. Fardo, F., Beck, B., Allen, M., & Finnerup, N.B. (2020). Beyond labeled lines: A population coding account of the thermal grill illusion, *Neuroscience & Biobehavioral Reviews*, 108, 472-479.
19. Ferrè, E. R., Iannetti, G. D., van Dijk, J. A., & Haggard, P. (2018). Ineffectiveness of tactile gating shows cortical basis of nociceptive signaling in the Thermal Grill Illusion. *Scientific reports*, *8*(1), 1-7.
20. Finlay, B.L. (2019). The neuroscience of vision and pain: evolution of two disciplines, *Philosophical Transactions of the Royal Society B*, 374*,* no. 1785: 20190292.
21. 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.
22. Flor, H. & Turk, D.C. (2013). Cognitive and Learning Aspects. In McMahon, S., Koltzenburg, M., Tracey, I., & Turk, D. (Eds.), *Wall and Melzack's Textbook of Pain* (pp. 256-272), Philadelphia: Elsevier Saunders.
23. Fodor, J.A. (1983). The Modularity of Mind, Cambridge, MA: MIT Press.
24. Fodor, J.A. (1985). Precis of the modularity of mind. *Behavioral and Brain Sciences*, 8(1), 1-42.
25. Fodor, J.A. (1989). Why should the mind be modular? In A. George (Ed.), *Reflections on Chomsky* (pp. 1-22). Oxford: Basil Blackwell.
26. Gibson, J.J. (1979). The ecological approach to visual perception. Boston: Houghton Mifflin.
27. Gligorov, N. (2017). Don't Worry, This Will Only Hurt a Bit: The Role of Expectation and Attention in Pain Intensity, *The Monist* 100(4), 501-513.
28. Green, E.J. (2021). The Perception-Cognition Border: A Case for Architectural Division, *The Philosophical Review,* 129(3), 323-93.
29. Gregory. R. L. (1970). The intelligent eye. London: Weidenfeld & Nicolson.

30. Gross, S. (2017). Cognitive Penetration and Attention. *Frontiers in Psychology*, 8, 221.
31. Hatfield, G. (2002). Perception as Unconscious Inference. In Heyer, D. & Mausfeld, R. (Eds.), *Perception and the Physical World* (pp.113-143). New York, NY: John Wiley & Sons.
32. Howhy, J. (2013). The Predictive Mind. Oxford: Oxford University Press.
33. Howhy, J., Roepstorff, A., & Friston, K. (2008). Predictive coding explains binocular rivalry: an epistemological overview. *Cognition,* 108(3), 687-701.
34. Jacobson, H. (2017). Pain and Cognitive Penetrability, In Corns, J. (Ed.), *The Routledge Handbook of Philosophy of Pain* (pp. 266-275). Abingdon: Routledge.
35. Klein, C. (2015). What the Body Commands: The Imperative Theory of Pain. Cambridge, MA: MIT Press.
36. Liberman, A. M., Cooper, F. S., Shankweiler, D. P., & Studdert-Kennedy, M. (1967). Perception of the speech code, *Psychological Review*, 74(6), 431-461.
37. Liberman, A. & Mattingly, I. (1985). The motor theory of speech perception revised, *Cognition*, 21(1), 1-36.
38. Linden, D.J. (2015). Touch. Viking Penguin: New York.
39. Longo et al. (2009). Visually Induced Analgesia: Seeing the Body reduces Pain, *The Journal of Neuroscience*, 29(39), 12125-12130.
40. Marr, D. (1982). Vision: A computational investigation into the human representation and processing of visual information. San Francisco: W.H. Freeman.
41. Martínez, M. (2011). Imperative Content and the Painfulness of Pain, *Phenomenology and the Cognitive Sciences* 10(1), 67-9.
42. Martínez, M. & C. Klein (2016). Pain Signals Are Predominantly Imperative, *Biology and Philosophy* 31(2), 283–98.
43. Mancini, F. et al. (2011). Visual distortion of body size modulates pain perception, *Psychological science* 22(3), 325-330.
44. Mandelbaum, E. (2018). Seeing and Conceptualizing: Modularity and the Shallow Contents of Perception, *Philosophy and Phenomenological Research,* 97(2), 267-83.
45. Mandelbaum, E. (2019). Modularist Explanations of Experience and Other Illusions, *Consciousness & Cognition*, 76, 102822.
46. Melzack, R. & Wall, P. (2008). The challenge of pain. New York: Penguin.
47. Melzack, R., Wall, P., & Ty, T.C. (1982). Acute Pain in an Emergency Clinic: Latency of Onset and Descriptor Patterns Related to Different Injuries, *Pain* 14(1), 33-43.
48. Mole, C. (2015). Attention and Cognitive Penetration. In Zeimbekis, J. & Raftopoulos, A. (Eds.), *The Cognitive Penetrability of Perception: New Philosophyical Perspectives* (pp. 218-238). Oxford: Oxford University Press.
49. Mylopoulos, M. (2021). The Modularity of the Motor System. *Philosophical Explorations,* 24(3), 376-393.
50. Nanay, B. (2018). Multimodal mental imagery. *Cortex,* 105, 125-136.
51. Ongaro, G. & Kaptchuk, T.J. (2019). Symptom perception, placebo effects, and the Bayesian brain, *Pain* 160(1), 1.
52. Orlandi, N. (2014). The Innocent Eye: Why Vision is Not a Cognitive Process. Oxford: OUP.
53. Palmer, S. E. (1999), Vision science: Photons to phenomenology. Cambridge, MA: MIT Press.
54. Pagnoni, G., & Porro, C.A. (2014). Cognitive modulation of pain and predictive coding: comment on "Facing the experience of pain: a neuropsychological perspective" by Fabbro & Crescentini, *Phys. Life Rev* 11, 555-557.
55. Pitcher, G.W. (1970). Pain Perception, *Philosophical Review* 79(3), 368-393.
56. Price, D.D., Finniss, D.G., & Benedetti, F. (2008). A Comprehensive Review of the Placebo Effect: Recent Advances and Current Thought, *Annual Review of Psychology*, 59, 565-590.

57. Prinz, J.J. (2006). Is the mind really modular? In Stainton, R.J. (Ed.), *Contemporary Debates in Cognitive Science* (pp. 22-36). Oxford: Blackwell.
58. Pylyshyn, Z. (1984). Computation and Cognition: Toward a Foundation for Cognitive Science. Cambridge: MIT Press.
59. Pylyshyn, Z. (1999). Is vision continuous with cognition?: The case for cognitive impenetrability of visual perception. *Behavioral and brain sciences*, *22*(3), 341-365.
60. Quilty-Dunn, J. (2020). Attention and Encapsulation. *Mind & Language*, 35(3), 335-349.
61. Ramachandran, V.S. (1988). Perception of shape from shading, *Nature*, 331(6152), 163-166.
62. Ratcliffe, R. & McKoon, G. (2008). The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks, *Neural Computation,* 20(4), 873-922.
63. Ramachandran, V. S., Rogers-Ramachandran, D., & Cobb, S. (1995). Touching the phantom limb. *Nature*, *377*(6549), 489-490.
64. Ramachandran, V.S. & Altschuler, E.L. (2009). The use of visual feedback, in particular mirror visual feedback in restoring brain function, *Brain* 132(7), 1693-1710.
65. Ramachandran, V.S., Brang, D., & McGeoch, P.D. (2009). Size reduction using Mirror Visual Feedback (MVF) reduces phantom pain, *Neurocase* 15(5), 357-360.
66. Rescorla, M. (2015a). Review of "The Innocent Eye: Why Vision is not a Cognitive Process" by Nico Orlandi. *Notre Dame Philosophical Review* (2015.01.10).
67. Rescorla, M. (2015b). Bayesian Perceptual Psychology. In. Matthen, M. (Ed.), *The Oxford Handbook of the Philosophy of Perception* (pp. 694-716). Oxford: Oxford University Press.
68. Ringkamp, M., Srinivasa, R.N., Campbell, J.N., & Meyer, R.A. (2013). Peripheral mechanisms of cutaneous nociception. In McMahon, S., Koltzenburg, M., Tracey, I., & Turk, D. (Eds.), *Wall and Melzack's Textbook of Pain* (pp. 1-30), Philadelphia: Elsevier Saunders.
69. Rock, I. (1975). An introduction to perception. New York: Macmillan.
70. Rock, I. & Victor, J. (1964). Vision and Touch: An experimentally created conflict between the two senses. *Science*, 143(3606), 594-6.
71. Sandkühler, J. (2013) Spinal cord plasticity and pain. In McMahon, S., Koltzenburg, M., Tracey, I., & Turk, D. (Eds.), *Wall and Melzack's Textbook of Pain* (pp.94-110), Philadelphia: Elsevier Saunders.
72. Schmidt R., Willis W. (2007). Encyclopedia of Pain. Springer: Berlin. https://doi.org/10.1007/978-3-540-29805-2_3363.
73. Scholl, B. J., & Leslie, A. M. (1999). Explaining the infant's object concept: Beyond the perception/cognition dichotomy. In Lepore, E. & Pylyshyn, Z. (Eds.), *What is cognitive science* (pp. 26-73). Oxford: Blackwell.
74. Schott, G.D. (2001). Delayed onset and resolution of pain: Some observations and implications, *Brain* 124(6), 1067-1076.
75. Senkowski, D., Höfle, M., & Engel, A.K. (2014). Crossmodal shaping of pain: a multisensory approach to nociception, *Trends in Cognitive Science*, 18(6), 319-327.
76. Seymour, B. & Dolan, R.J. (2013). Emotion, Motivation, and Pain. In McMahon, S., Koltzenburg, M., Tracey, I., & Turk, D. (Eds.), *Wall and Melzack's Textbook of Pain* (pp.248-255), Philadelphia: Elsevier Saunders.
77. Shevlin, H. and Friesen, P. (2021). Pain, placebo, and cognitive penetration, *Mind & Language*, 36(5), 771-797.
78. Sun, J., & Perona, P. (1998). Where is the sun? *Nature Neuroscience,* 1(3), 183–184.
79. Thunberg, T. (1896). Förnimmelserna vid till samma ställe lokaliserad, samtidigt pågående köld-och värmeretning, *Uppsala Läkfören Förh* 2(1), 489-495.
80. Tabor, A., Thacker, M. A., Moseley, G. L., & Körding, K. P. (2017). Pain: a statistical account. *PLoS computational biology*, *13*(1), e1005142.

81. Trommershäuser, J., Kording, K., & Landy, M.S. (2011). Sensory Cue Integration. Oxford: Oxford University Press.
82. Wager, T.D. & Fields, H.L. (2013). Placebo Analgesia. In McMahon, S., Koltzenburg, M., Tracey, I., & Turk, D. (Eds.), *Wall and Melzack's Textbook of Pain* (pp.362-373), Philadelphia: Elsevier Saunders.
83. Wall, P.D. (1979). On the Relation of Injury to Pain, *Pain* 6(3), 253-264.
84. Wiech, K., Vandekerckhove, J., Zaman, J., Tuerlinckx, F., Vlaeyen, J.W.S., & Tracey, I. (2014). Influence of prior information on pain involves biased perceptual decision-making, *Current Biology* 24(15), R679-R681.
85. Wiech, K. (2016). Deconstructing the sensation of pain: The influence of cognitive processes on pain perception, *Science* 354(6312), 584-587.
86. Wittkopf, P.G., & Johnson, M.I. (2017). Mirror therapy: A potential intervention for pain management, *Revista da Associação Médica Brasileira* 63(11), 1000-1005.
87. Wu, W. (2014). Shaking Up the Mind's Ground Floor: The Cognitive Penetration of Visual Attention, *Journal of Philosophy*, 114 (1), 5-32.