The Pain System is Not a Bodily Disturbance Detector

Forthcoming in Ana Cuevas-Badallo, Mariano Martín-Villuendas & Juan Gefaell (eds.), *Life and Mind: Theoretical and Applied Issues in Contemporary Philosophy of Biology and Cognitive Sciences.* Springer.

This is an accepted version of the manuscript. Please cite the published version.

Tiina Rosenqvist
Dartmouth College, Hanover, NH, USA

Abstract What is the function of pain? A popular view in contemporary philosophy is that the pain system is a bodily disturbance detector: pain states track/detect and represent bodily disturbances and the phenomenal character of the (sensory dimension of) pain supervenes on this representational content. The view can accommodate paradigmatic pain cases, e.g., when pain follows from stepping on a nail. Once we consider more complex pain phenomena, however, it has seemingly little to offer. In this chapter, I discuss dissociation between pains and bodily disturbances, variation in pain thresholds, the effects of repeated stimulation on experienced pain intensity, and the modulation of pain experience by contextual factors. I argue that these phenomena suggest that the pain system is not a bodily disturbance detector, but a sophisticated security system.

Keywords Pain; Function; Perception; Action guidance; Detection

1 Introduction

Though most of us are intimately familiar with pain, giving a coherent philosophical account of pain is difficult. Our intuitions conflict, the referent of our pain talk is unclear, and there seem to exist plausible counterexamples to virtually every account. In paradigmatic pain cases, the pain experience is caused by a physical disturbance to some specific part of the body, and many philosophers maintain that pain experiences detect/track and represent such disturbances (e.g., Tye 1995a, 1995b; Dretske 1997; Hill 2009; Cutter and Tye 2011). A natural companion to such views is the conceptualization of the pain system as a sort of a bodily disturbance detector, the primary function of which is to inform us about the damaged condition of our bodies.

I have two main objectives in this work. First, I argue that, because the explanatory power of the bodily disturbance detector view of pain is severely limited, we ought to reject the view. Second, I motivate an alternative conception of the function of the system: the pain system is a sophisticated security system, the primary function of which is to keep us safe. I proceed as follows: I first motivate and explicate the bodily disturbance detector view (Sect. 2) and consider some previous criticisms (Sect. 3). I then make my own case against the view. I argue that the view has

trouble accommodating a wide variety of pain phenomena: dissociation between pains and bodily disturbances, variation in pain thresholds, effects of repeated stimulation on experienced pain intensity, and contextual modulation of pain experience (Sect. 4). I then introduce the sophisticated security system view of pain as an alternative to the detector view and explain how it can accommodate and explain the phenomena just mentioned (Sect. 5). My exposition of the security system view is programmatic, but I also consider some questions that immediately arise and discuss ways in which the view could be developed further and applied to clinical and research settings (Sect. 6). Finally, I conclude that we have good reason to accept the security system view, or something close to it (Sect. 7).

2 The bodily disturbance detector view of pain

Suppose that you step on a nail and feel a throbbing pain in your foot. This is a paradigmatic pain scenario where the injury causes your pain experience, and you experience the pain in the injured body part. Your pain experience is clearly body-directed, perhaps even "transparent" in the sense that you experience the qualities of the pain as qualities of your foot, not as intrinsic features of your experience itself (Tye 2002). If your pain experience is directed at your foot (*i.e.,* it is "about" your foot), we can say that it has content. It is often assumed that this content is *representational*: your pain experience represents your foot being a certain way.

If pain experience has representational content, then how is that content determined? There are several plausible stories to be told here. For example, it might be that the content of your experience is determined by the function of the pain system: the system represents *some property* or *state P* if and only if it has the function of informing about *Ps* (Dretske 1997, p. 2). This story is teleological: pain representations are produced by a mechanism selected for the purpose of informing about *Ps*. On the other hand, it might be that your pain experiences "track" *Ps.* On this alternative view, pain experiences represent *Ps* if and only if they causally co-vary with *Ps* under typical/optimal conditions (Tye 1995a, p. 113; Cutter and Tye 2011, p. 91). In this tracking story, no explicit reference is made to the function of the system, though it seems natural to understand tracking as something that the system has the function to do. After all, the tracking theorist needs an account of why pain experiences track *Ps,* and many philosophers find it natural to answer that it is the function of the pain system to do so. As Casser writes, if the function of pain were *not* to inform about *P*, "it would be surprising to find that pain still covaries with [*P*] in the relevant sense, even though pain is not in the business of tracking it"

(2021 p. 367).[1] For the remainder of this chapter, I will follow Casser in assuming that both the Dretskean view and the tracking view are best understood as involving a commitment to the claim that the function of the pain system is to detect or track (and inform about) *Ps*.[2]

It also *feels* a certain way for you to undergo a pain experience. It is an attractive idea that the phenomenal character of the (sensory dimension of) pain supervenes on the representational content of pain, *i.e.,* that there can be no difference in the phenomenal character without a difference in the representational content (Byrne 2001; Cutter and Tye 2011).[3] Such "representationalism" about pain experience makes the phenomenal seem less mysterious and more amenable to naturalistic treatment. A strong version of this claim simply equates the phenomenal character with a representational content of a certain kind (Tye 2002; Dretske 1997).

If the representationalist account of pain experience is correct, then what is *P*? Philosophers differ somewhat in their answers, but the basic idea is that pain experiences represent tissue damage or some other mind-independent disturbance to some part of the body. For example, Tye proposes that "pains are sensory representations of *bodily damage* or *disorder*…they are mechanical responses to the relevant bodily changes in the same way that basic visual sensations are mechanical responses to proximate visual stimuli" (1995a, p. 113; italics mine). The phenomenology of pain includes certain sensory properties, presumably at least a specific location, a specific intensity, and a specific quality (Hill 2017, p. 63). You might have a severe throbbing pain in your foot, or a mild ache in your knee, for example. If the phenomenal character of your experience supervenes on the

[1] If the notion of representation is to be a useful notion, it needs to leave space for the possibility of misrepresentation. And what is misrepresentation, if not the failure of the system to do what it is supposed to do (in typical/optimal conditions)? In Dretske's words, "there is information without functions, but there is no representation without functions" (1997, p. 4; see also Hatfield 1990, p. 264). Note that Cutter and Tye, in discussing the tracking story, write that we are in optimal conditions "just in case we are in conditions of the sort that our [relevant perceptual system] was designed to operate in by natural selection or by analogous processes in the course of ontogenic development" (2011, p. 19). To say that the pain system was designed to operate (*i.e.,* track the relevant states) in optimal conditions comes very close to saying that it is the function of the system to track those states under such conditions. In fact, later in the same paper the authors explicitly state that it is "overwhelmingly plausible that the internal states we undergo when we experience pain have the biological function of indicating that some bodily disturbance is apt to harm us" (ibid., p. 108, n11).

[2] For ease of exposition, I shall be assuming a broadly etiological (Wright 1973, 1976) conceptualization of biological functions. Readers sympathetic to alternative conceptualizations (e.g., Cummins 1975) are free to modify my argument to better fit their preferred view. In fact, as I briefly discuss in Sect. 4.2, my argument challenges the *general* idea that the function of the pain system is to represent bodily disturbances, even if such function is altogether divorced from biological function.

[3] I focus on the sensory aspect of pain and bracket the issue of whether pain experiences have other kinds of contents. For example, some philosophers maintain that pain states have evaluative contents (e.g., Bain 2013) which account for the "badness" of pain. Sometimes the evaluative and sensory contents are integrated, as in Cutter and Tye's proposal that the content of a pain experience is something like "there is a bodily disturbance of (physiological) type d in location l, and d is bad for A to degree x" where A is the subject of the experience and the property of being bad for A is analyzed as being apt to harm A (2011, p. 99).

representational content and if the representational content is that there is a (mind-independent) disturbance with certain features in your right foot, then the represented features are supposedly some mind-independent features of the disturbance (Pautz 2014, p. 248; see also Cutter 2017, p. 34). In what follows, I will use 'bodily disturbance' as an umbrella term intended to capture the idea that the relevant state is some mind-independent state related to potential or actual damage or injury.

Note that if the content of pain experience is mind-independent and determined either by the function of the system or through a tracking relation, then pain experiences have accuracy conditions. A particular pain experience is accurate if there exists a disturbance of the appropriate sort in the relevant bodily location (e.g., a puncture wound in your foot), and inaccurate if there doesn't. In other words, pain experiences can *mis*represent, just as, *say*, visual experiences can misrepresent. Some pains are illusory (e.g., when a mildly noxious stimulus is perceived as excruciatingly painful) and others are hallucinations (e.g., some phantom limb pains).[4]

Putting together these different strands (tracking/Dretskean view of content determination, representationalism about phenomenal pain experience, and physical content), we arrive at the bodily disturbance detector view of pain:

> [BDD] The pain system is a bodily disturbance detector: the function of pain states is to track/detect and represent physical, mind-independent bodily disturbances and their mind-independent properties, and the phenomenal character of (the sensory dimension of) pain experience supervenes on this representational content.

The detector metaphor has played an influential role in philosophy of perception; an analogy is often drawn between sense perception and human-made instruments designed to detect and represent some physical states of the world.[5] The metaphor is simple, effective, and intuitive. It accords well with popular naturalistic theories of mental representation and with the ordinary conception of perception as providing a transparent window into the world. It also provides a unified picture of all sensory perception: our perceptual systems are in the business of accurately representing the

[4] See, e.g., Hill 2017, p. 63. Tye, too, accepts that pain experiences can *mis*represent, even when they are genuine. For example, the referred pain of a heart attack (felt in one's left arm) is a genuine pain that misrepresents the location of the relevant disturbance. Tye writes: "A man who reports to his doctor that he has a pain in his left arm is not taken to have lied, if it is discovered that the real cause of his pain lies in his heart. Such a man has a pain in his left arm, but in this case he *is* under a kind of illusion: there really is nothing *in his left arm* which is hurting him" (1995b, p. 228).

[5] Dretske famously uses instruments to explicate his account of mental representation. For example, immediately after expressing his "fundamental idea" (*i.e.,* that a system represents a property *iff* it has the function of informing about that property), he describes the operations of a speedometer to illustrate what he means (1997, p. 2).

mind-independent world around us or, in the case of pain perception, in the business of accurately representing disturbed states of our own body parts. Despite these advantages, I will argue that the detector metaphor is seriously misleading when it comes to pain.[6]

I focus on BDD because it is arguably the simplest and most natural formulation of the idea that the pain system is a bodily disturbance detector. This is because BDD doesn't attribute systematic error to pain states but allows that the contents of pain experiences could be entirely veridical. More complex views state that pain experiences systematically misrepresent bodily disturbances, e.g., by attributing to those disturbances properties that they do not actually have (Hill 2017; see also Hill 2009, Ch. 6). What I say in this chapter has relevance, with some modifications, to any view of pain which takes pain experience to be primarily the perception of the current states of one's body but for the sake of simplicity I use BDD as my primary target.[7]

3 Challenges to BDD

A variety of objections have been raised against BDD and the views in its vicinity. Some have asked why—if feeling pain is perceptually representing a disturbance to some part of the body—pain beliefs ostensibly report on the *experience* and not on the disturbance itself (Aydede 2009, 2017). I can believe that there is pain in my foot and at the same time believe that there is nothing physically wrong with my foot. There is no inconsistency here; both beliefs can be simultaneously correct. A proponent of BDD might respond that we use the term 'pain' to refer to pain experience simply because "pain bothers us, often quite terribly" (Hill 2017, p. 67). In other words, we *care* about our pain *experiences*, even in the absence of bodily disturbances, and our language use reflects what we care about. In the end, the debate comes down to what we take to be the defining features of perceptual experience. For Aydede, what matters are the types of beliefs and judgments that the

[6] Others have also found the detector metaphor misleading. Arguments have been raised against the detector views of specific perceptual systems and against the detector view of perception in general. For example, Hoffman et al. (2015) use evolutionary games and genetic algorithms to empirically investigate the idea that natural selection favors veridical perceptions (and thus produces "detectors" of various sorts). Their conclusions are straightforward: "Natural selection discards veridical perceptions and promotes interface strategies tuned to fitness" (ibid., p. 1488). Such "interface strategies" are strategies that present the organism's environment in a way that is useful to that organism, e.g., by guiding successful action.

[7] BDD is a proper subset of a set of theories that Casser (2021) groups under the label "information-gathering story of pain" (IGS). According to IGS, the primary function of pain is to inform organisms about damage to their bodies. Such information gathering need not involve representations or—if representations are involved—do the contents of those representations need to be fully externalist/physical. That said, the simplest version of the information-gathering story is that pain experiences represent physical states of the body and that the phenomenal content of pain experience supervenes on the representational content.

experiences directly ground. He worries that there is no good reason to take pain experiences to be genuinely perceptual if pain-attributing beliefs and judgments (e.g., "I have a throbbing pain in my foot") are never *directly* (*de re*) about the relevant kinds of disturbances being instantiated in the relevant body parts (in Aydede's view this is evidenced by the fact that we can take such beliefs and judgments to be true even in the absence of disturbances), when paradigmatic perceptual experiences (e.g., seeing a cup as round) typically *do* give rise to *de re* judgments about their extramental objects as such (e.g., "this thing is round"). In response, those who take pain experiences to be genuinely perceptual might emphasize, for example, the similarities in phenomenology between pain experiences and paradigmatic perceptual experiences, such as the observation that pain experiences typically assign a bodily location to pains (Hill 2017, p. 63).

Related to the first objection, others have argued that, when it comes to pain, *esse est percipi—i.e.,* that one cannot have a pain without having a pain experience. Because BDD and its theoretical cousins identify pains with mind-independent bodily disturbances (at least on the assumption that pain experiences represent pains), and because bodily disturbances can occur in the absence of pain experiences, these views are inconsistent with the *percipi* intuition (Pautz 2010). But perhaps our intuitions are simply confused. It seems that we apply the concept of pain both to the physical disturbances (pains-as-objects) and to the experiences those disturbances tend to cause (pains-as-experiences). If we get clear on this distinction, we can make sense of the cases where the two come apart (Byrne 2001, p. 229; see also Hill 2009, Ch. 6). In the case of some phantom limb pains, for example, one might have pain-as-experience, but not pain-as-object. If a pain is felt in a phantom hand, the hand isn't there and so the relevant sort of bodily disturbance (*i.e.,* pain-as-object) isn't there either. The experience, which is one type of pain, misrepresents a bodily state, which is another type of pain.[8] Then, plausibly, one could also have pain-as-object, but not pain-as-experience, e.g., in the case of injuries that are not experienced as painful. Still, as Pautz point out, we *never* seem to apply the concept of pain to bodily disturbances in the absence of pain experience. For example, there is something deeply counterintuitive about the idea that cadavers could have pains (2014, p. 240). The body parts of cadavers can have various properties (shapes,

[8] It could be argued that at least in some cases of phantom limb pain (PLP), there does exist pain-as-object, but the location of the pain-as-object is misperceived. The physiology of PLP remains poorly understood but according to one theory, it is caused by the irritation of the severed and damaged nerve endings in the residual limb (see Hanyu-Deutmeyer et al. 2022). If this were true, then the PLP experience could be said to represent the nerve irritation but misrepresent its location, at least if peripheral nerves are understood as part of the body in the relevant sense, *i.e.,* if they are taken to be the relevant kind of bodily tissue whose disturbed condition our pain experiences represent (this is not a given). That said, the causes of PLP are complex and thought to involve many peripheral and central nervous system factors. The primary cause of some cases of PLP might be the sensitization of the central nervous system, in which case we would have pain-as-experience in the absence of pain-as-object. And, in any case, there are other putative cases where the two types of pains come apart, such as fibromyalgia (e.g., Julien et al. 2005).

sizes, etc.) and conditions (damage, etc.), but intuitively not pains. For Pautz this is reason enough to conclude that pains aren't mind-independent.

Finally, philosophers have made empirical cases against BDD-type views. Casser (2021) argues that neither biological nor neurophysiological evidence supports the conception that the function of pain is to inform about bodily disturbances (see also Klein 2015). Pautz (2010, 2014) takes psychophysical and neuroscientific data to reveal two things: (i) that there is a structural mismatch between the intensity of pain and the physical magnitudes of the corresponding bodily disturbances and (ii) that pain intensity depends on internal (neural) states. The case I make against BDD has affinities with both Casser's and Pautz's approaches. Like Casser, I focus on the function of the pain system. I seek to show that BDD does a remarkably poor job at accommodating and explaining a variety of pain phenomena and should therefore be rejected. Unlike Casser, I contrast BDD with a different conception of the function of pain, which I suggest better explains the empirical data. In other words, I advance a positive view on the function of the pain system. Like Pautz, I argue that psychophysical evidence (in my case supplemented with phenomenological observation) puts pressure on certain kinds of views about pain. But whereas Pautz argues against phenomenal externalism about pain (*i.e.,* the view that pains are mind-independent bodily states), I argue against the idea that the pain system is a bodily disturbance detector. When it comes to deciphering the mysteries of pain, Pautz thinks that our best bet is to adopt a combination of primitivism about pains and projectivism about pain experience. I think that our best bet is to begin with a different conception of the function of the pain system and to build up from there. In the present work, I focus *only* on the function question and do not defend a positive view on the nature of pain itself. [9]

4 Complex pain phenomena: a case against BDD

To support my claim that the explanatory power of BDD is severely limited, I focus on four kinds of pain phenomena: dissociation between pains and bodily disturbances, interpersonal and intrapersonal variation in pain thresholds, temporal adaptation to repeated noxious stimulation, and contextual modulation of pain experience. I propose that each of the four phenomena puts pressure on BDD and that together they make a strong cumulative case against the idea that the goal of the pain system is to detect bodily disturbances.

[9] In Pautz's view, pains are "wholly chimerical" and exist nowhere in the world (2014, p. 295). Others who are suspicious of the idea that pains are bodily states (even *experienced* bodily states) might argue that they are subjective mental states instead. Readers sympathetic to either of these views could read me as providing a conception of the function of the pain system that complements their preferred view of the nature of pain (many thanks to an anonymous reviewer for pointing this out). I will revisit the issue of the nature of pain in Sect. 6.1.

4.1 The dissociation of pains and bodily disturbances

Bodily disturbances, even severe ones, can occur without pain experiences. Consider the case of Rodney Fox being attacked by a great white shark during a 1963 spearfishing competition in Australia. In a *Guardian* article, Fox describes getting brutally dragged and mauled by the shark yet claims to have only experienced pain after he had been rescued from the water. "It was only when I was lying down on the boat that the waves of pain arrived," he is quoted as saying.[10] Other examples abound: soldiers wounded in World War II, people engaging in body suspension (a practice where a person's skin is pierced with sterile metal hooks from which the person is then hung), and victims of car crashes and other accidents. In all these cases, bodily disturbances are undeniably present, but people report feeling no pain.[11] [12]

On the other hand, pain experiences can also occur in the absence of the relevant kinds of bodily disturbances. Lots of people suffer from medically unexplained pain, where pain is felt in a body part where no evidence of a physical disturbance is found. One example is fibromyalgia, a condition characterized by generalized musculoskeletal pain, often in the absence of observable musculoskeletal problems. Other examples include some neuropathic pains as well as many types of back pains and headaches. Phantom limb pain shows that pain can even be felt in a body part that is no longer there.

It has been suggested that BDD-like views cannot account for this dissociation between pains and bodily disturbances,[13] but this conclusion is not obvious. When pain experience occurs without a bodily disturbance, the experience could be said to

[10] Ammar Kalia, "'I saw a big set of white teeth coming towards me': the people who survived terrifying wild animal attacks." Guardian 4 Jan 2022: https://www.theguardian.com/environment/2022/jan/04/saw-big-set-of-white-teeth-coming-towards-me-people-who-survived-wild-animal-attacks (accessed May 27, 2023).

[11] See Beecher (1956) on wounded soldiers and Melzack et al. (1982) on emergency room patients. My claim about the painlessness of body suspension relies on reports from people who routinely engage in this practice.

[12] Note that some of these cases might involve a phenomenon analogous to that of inattentional blindness where fully visible (but unexpected) objects are not noticed or consciously experienced because attention is engaged elsewhere (Mack and Rock 1998). Evidence suggests that we can also fail to notice noxious stimuli (such as the sound of fingernails scratching a chalkboard) when engaged in a task that focuses our attention elsewhere (Wayand et al. 2005). If the dissociation between pains and bodily disturbances in particular cases is explained by something like inattentional blindness, then such cases arguably tell us very little about the goal of the pain system (at least on the assumption that inattentional blindness tells us very little about the goal of the visual system). That said, attentional effects cannot explain *all* the cases where bodily disturbances occur without accompanying pain experiences; people often report feeling no pain *even when they specifically attend to the disturbed condition of their bodies*. For example, people engaging in body suspension are often curious about their pain level and carefully attend to it. Beecher's soldiers, all of whom had extensive wounds and who were "clear mentally" and not in shock, were asked the question, "As you lie there, do you feel any pain?" Many reported feeling no pain, even though they would complain "in a normal manner at rough handling of their wounds, or at inept venipunctures" (1956, pp. 1609–1611).

[13] See, e.g., Klein 2015, Ch. 3. Klein is mainly drawing from Wall (1979) here.

simply *mis*represent: the content of the experience would be something akin to "there exists a disturbance (with such-and-such properties) in this bodily location," but the content would be false (see, e.g., Hill 2017, p. 63). In some ways, this response is intuitive. After all, *something* has surely gone wrong in the case of fibromyalgia pain or back pain that presents without tissue damage. Such pains are usually maladaptive; they are more harmful than helpful (fibromyalgia pain might cause people to refrain from health-promoting physical exercise, for example). It is not clear, however, that BDD provides the best explanation for these cases. In addition, the sheer commonness of the cases poses a problem. Under BDD, pain in the absence of relevant tissue damage is analogous to visual hallucination; the experience represents something that isn't there. If BDD entails that such "pain hallucinations" are commonplace, it owes us an explanation as to why pain perception is more prone to these errors than other perceptual modalities.[14] On the assumption that all sensory systems are detectors, what explains the high number of false positives in the case of pain?

When it comes to painless injuries, it could be argued that the conditions in which they occur are simply too far removed from the optimal or typical conditions in which our pain systems can be expected to function normally. If our visual systems cannot operate in the dark, then perhaps our pain systems cannot operate properly when we are being mauled by wild animals or suspended by hooks placed under our skin. At least in the former case we might be dealing with stress-induced analgesia (SIA), *i.e.,* an inbuilt pain suppression response to stressful or fearful stimuli. SIA likely serves an important adaptive purpose: it allows us to focus on surviving a dangerous situation instead of tending to our injuries (e.g., Butler and Finn 2019, p. 185). If you wish to survive a shark attack, you ought to focus on getting away!

Does SIA help the proponent of BDD account for painless injuries? Perhaps she could argue that the typical or optimal conditions for detecting bodily disturbances do not include high-stress situations where pain experience would likely be maladaptive, and that SIA results from a mechanism that *inhibits the normal and proper functioning of the pain system*. At first, this seems plausible. After all, the pain system involves the activity of peripheral nociceptors, *i.e.*, specialized afferent fibers in the body tissues that are activated by noxious stimuli of different kinds: thermal, chemical, and mechanical (Dubin and Patapoutian 2010).[15] The nociceptors

[14] It is difficult to know exactly how prevalent pain experience in the absence of bodily disturbance is, but e.g., fibromyalgia, described as "the diminished quality of life related to generalized body pains and physical and psychological symptoms *that occurs in the absence of a clear pathologic cause*" (my italics), is fairly common; some 1.75% of adults (18 years or older) in the US satisfy the diagnostic criteria (Walitt et al. 2015).

[15] Nociceptors are categorized according to the properties of their axons into polymodal nociceptors (associated with unmyelinated C-fibers) and two types of Aδ nociceptors (associated with thinly myelinated Aδ-fibers): mechanosensitive and mechanothermal (Purves et al. 2004, p. 209). When stimulated simultaneously, the polymodal nociceptors and Aδ nociceptors give rise to different kinds of pain. Aδ stimulation leads to a sharp, brief, and well localized "first pain" whereas polymodal nociceptor stimulation leads to a dull and poorly localized

project to the dorsal horn of the spinal cord and from there various types of ascending pathways carry nociceptive information to the brain. Importantly, however, the transmission of the signals is modulated by descending pathways from the brain to the spinal cord (Purves et al. 2004). In the case of SIA, nociceptive signals are *inhibited* by such pathways. This means that SIA could perhaps be conceptualized as a kind of an emergency kill switch for the pain system.[16]

The problem with this response is that the inhibitory mechanisms involved in SIA are generally considered to be part of the pain system itself.[17] And so it looks like the pain system *is* operating in these conditions, even if its operation results in the suppression of pain experience. This makes the case non-analogous with the visual system case. In the dark, there aren't enough inputs for the visual system to inform us about our surroundings. In the painless injury case, however, the nociceptive inputs from the affected body parts are *actively suppressed*. As Casser points out, correctly, in my opinion, the challenge here is to "explain why a system would purposefully prevent system-relevant information from transmission if it is the system's function to perform such transmissions" (2021 p. 371).[18] In other words, if the goal of the system is to detect/track and represent bodily disturbances, then why would the system itself actively work against this goal? Once again, BDD owes us an explanation.

4.2 Interpersonal and intrapersonal variation in pain thresholds

Even if BDD could account for painless injuries and pain in the absence of injuries, other problems remain. For example, there exists considerable interpersonal and intrapersonal variation in pain thresholds, even among healthy subjects in perfectly ordinary conditions. Experimental data suggests that pain thresholds increase with

---

"second pain" that normally outlasts the noxious stimulus presentation (Price and Barrell 2012, Sect. 7.4). If the phenomenal character of pain experience supervenes on the content of the experience, then the content of first pain and second pain must be different. This is something that the proponent of BDD should also account for.

[16] Hardcastle suggests that instead of one pain system we have two separate systems: a nociceptive system which "keeps us informed of the status of our bodies" and a top-down inhibitory system which "shuts down the [first system] when flight or fleeing is immanent" (1999, pp. 130-134). Hardcastle does not, however, advocate BDD or the kill switch view. For her, nociception does not equal pain perception, and the inhibitory system is *part* of pain processing.

[17] Melzack and Wall's (1965) gate control theory revolutionized our understanding of pain: the idea is that there are gates in the spinal cord that can block nociceptor signals or let them through. Though the original theory is now considered to be oversimplified, the understanding of pain processing as essentially involving the modulation of nociceptor outputs in the spinal cord has remained. As Butler and Finn write, "It is now clear that pain is processed through multiple, interweaving receptor-mediated pathways utilising the excitatory and inhibitory amino acids, monoaminergic, opioid, and endocannabinoid systems among others" (2019, p. 196).

[18] Casser (2021) also considers a stronger version of this response which appeals to the filtering out of irrelevant information in other sensory modalities such as taste (we don't taste our own saliva). He responds that whereas this sort of filtering of information can be understood as a form of information control, the same cannot be said about pain. In pain modulation, information *relevant* to the functioning of the system is blocked or reduced.

age (Lautenbacher et al. 2017; see also Goksan et al. 2015), and that there are significant gender differences in experimentally induced pain (Paulson et al. 1998; Bartley and Fillingim 2013). This means that stimuli that used to be experienced as painful often cease to give rise to pain experiences as individuals get older, and when two healthy young subjects are exposed to the same noxious stimulus, only one might experience the stimulus as painful. Imagine that two subjects, A and B, are administered a thermal stimulus of 43°C.[19] A rates the stimulus as painful, B doesn't. On the assumption that the ratings faithfully reflect the subjects' experiences, which experience is accurate: the painful or the painless one? Whose pain system is working properly, and whose isn't?

In some cases, variation in pain thresholds can be explained by deterioration. For example, the heightened pain thresholds in older subjects likely reflect an age-related decline in perceptual ability, comparable to a decline in visual acuity. In addition, interpersonal variation in absolute thresholds (*i.e.,* in the smallest levels of stimulus that can be detected 50% of the time) is very common in general,[20] and it isn't always clear when a stimulus *should* be detectable, unless we take the lowest absolute threshold measured as our standard. We also do not normally consider this kind of variation a problem for theories of other perceptual modalities, so why should it be a problem for BDD? Perceptual sensitivities differ and that's nothing new.

But there *is* a relevant difference between pain experiences and experiences in other perceptual modalities that the variation in pain thresholds highlights. With vision, for example, we have standardized tests that can be used to determine the smallest letters you can read on a chart. We know what the letters are, and we know when a person's visual experience of the letters is *accurate*. We don't have *these* kinds of standardized tests for pain acuity.

But let us dig a little deeper. The visual acuity test was designed for the purpose of measuring how well a person can see from a distance. This ability is required for several important tasks, and when one's visual acuity dips below a certain threshold, it can be easily corrected with glasses, contacts, or even surgery. That we do not have a comparable test for pain thresholds might just reflect the fact that we have no need for one; perhaps it doesn't matter very much whether a 43°C thermal stimulus feels painful or just hot.

[19] For the sake of the argument, we can even stipulate that the thermal stimulus gives rise to identical nociceptor firing patterns in A and B.

[20] For example, Humes et al. (2009) report age-related declines in threshold sensitivity in several sensory modalities, including vision, hearing, and touch.

Nevertheless, *if* it is assumed that perception aims at accurately representing mind-independent properties of, or events occurring in, the environment (*i.e.*, if BDD is generalized into a theory of perception as a whole), experiences in paradigmatic perceptual modalities generally allow for the *comparison* of perceivers' sensitivities. For example, a child who hears a high-frequency sound would be said to detect something that *is* there, whereas the child's grandfather would be said to fail at this task should his auditory experience not represent the stimulus at all. With pain, the situation is different. We can probe people with thermal stimulators all we want, but doing so will not help us determine the temperature at which a stimulus becomes *genuinely* noxious and something detectable by the pain system.[21] So, in the end, it is not just that we do not know whose pain experience is accurate, *we also do not know how the question should be answered in the first place*.

A big problem here is the lack of a clear and satisfactory definition of 'bodily disturbance.' An even bigger problem is the likely possibility that a clear and satisfactory definition will never be found because pain states do not neatly map onto any physical states (see Klein 2015, pp. 36-42). This erodes the motivation for the view that the pain system functions to detect such states. And while similar arguments have been raised against philosophical views which identify colors with physical, mind-independent properties of the external world (see, e.g., Hardin 1993; Cohen 2009), in the case of color it is generally clear what the relevant physical property is supposed to be (surface spectral reflectance, etc.). With BDD, this is lacking. The pain system is conceptualized as a detector, without a clear answer to the question of what it is supposed to detect.[22]

Because the notion of bodily disturbance is so vague, the worry is that BDD cannot tell us when the pain system is functioning properly. A focus on interpersonal variation in pain thresholds helps make this problem apparent. The problem is magnified by the fact that there is also *intra*personal variation in pain thresholds that isn't aging-related. Evidence suggests that human subjects can train themselves to suppress pain and increase their pain thresholds. For example, functional Magnetic Resonance Imaging (fMRI) has been used to teach people to activate the pain-controlling parts of their brain at will (deCharms et al. 2005). Conscious control over pain, if genuinely possible, tells us that pain thresholds are variable even in the

[21] Although pain scales and other tools can be employed to test subjects' sensitivity to changes in noxious stimulation, such tests cannot tell us whether a stimulus should be experienced as painful in borderline cases. For additional discussion of pain scales, see n23.

[22] It could be argued that there *is* an answer to the question of when a stimulus becomes *genuinely* noxious (*i.e.,* a bodily disturbance of the appropriate kind) but that we lack epistemic access to that answer. Analogous debates occur in philosophy of color where specifying the "true" fine-grained colors of objects has been shown to be equally difficult and where some proponents of physicalist views of color have responded that certain color facts might be unknowable (e.g., Byrne and Hilbert 2003). But, again, with color we have a physical property that competing chromatic experiences all supposedly represent (some more accurately than others, or so the story goes). With pain, we do not even know what the relevant property is supposed to be.

short term. This raises the question: why would a system whose goal is to detect mind-independent bodily states exhibit such plasticity in its operation? It might well be that the changes in sensitivity are adaptive (the way SIA is thought to be adaptive), but the problem for BDD is that such *adaptive plasticity* does not seem compatible with the idea that pain systems are in the business of mechanistically detecting mind-independent bodily states.

Now, it could be argued that veridical/accurate representation and fitness do not always go hand in hand. Pain systems could have a separate biological function in addition to their representational function. If so, then the *effects* of representational states could be adaptive even when the states themselves are in error (see, e.g., Burge 2010, Ch. 8), and adaptive plasticity would not have to be in tension with the representational function. The problem with this response is that, ideally, we would want an explanation of how pain states come to have the contents that they have. BDD theorists generally acknowledge this, which is why they explain content determination in terms of a tracking relation (in design conditions) or a biological function of the system. And even if representation were taken to be primitive, *we would still need a good reason for thinking that the pain system has a representational function in the first place*. I argue that complex pain phenomena speak against the view that pain states represent mind-independent bodily disturbances. If my arguments succeed, then separating representational function from biological function does little to help BDD.

4.3 Temporal adaptation: habituation and sensitization

In the previous two sections, I considered the dissociation of pains and bodily disturbances in unusual circumstances, and the difficulty of determining when pain should be felt in perfectly ordinary situations where the stimulus is only mildly noxious, if at all. In this section, I consider the dissociation of the severity of bodily disturbances and the intensity of accompanying pain experiences.

According to BDD, when the pain system is functioning properly, pain experiences optimally track/detect and represent mind-independent bodily disturbances and their mind-independent properties. I mentioned above that the phenomenology of pain includes at least three sensory properties: a specific location (e.g., the sole of your foot), a specific intensity (e.g., severe), and a specific quality (e.g., throbbing). Consider Tye's view on what different sorts of pain experiences represent:

> A twinge of pain represents a mild, brief disturbance. A throbbing pain represents a rapidly pulsing disturbance. Aches represent disorders that occur *inside* the body rather than on the outside (…). A

> stabbing pain is one that represents sudden damage over a particular well-defined bodily region (…). In each of the above cases, the subject of a pain undergoes a sensory representation of a certain sort of bodily disturbance. The disturbances vary with the pains (1995a, 113).

Pautz (2014, p. 249) notes that pain experience exhibits *response expansion:* the relationship between pain intensity and the severity of bodily disturbances is described by a power function with an exponent greater than 1. This means, for example, that a small increase in stimulus temperature (e.g., from 47°C to 50°C) can lead to a large increase in experienced pain intensity, and that the relationship between pain intensity and the physical magnitudes of bodily disturbances is somewhat "messy." Another phenomenon illustrating the messiness of this relationship is *temporal adaptation* to repeated noxious stimulation. Both phenomena, but especially temporal adaptation, give reason to think that BDD cannot account for the way in which experienced pain intensity varies.

Temporal adaptation is a well-known psychophysical fact. For example, Jepma et al. (2014) report that adaptation effects due to repeated stimulation "strongly modulate" the experienced painfulness of thermal stimuli. On average, repeated stimulation of the same skin site produces habituation (decrease in pain intensity) for all but the highest stimulus temperatures, whereas repeated stimulation across different skin sites produces sensitization (increase in pain intensity).[23] These patterns of adaptation clearly show that the experienced pain intensity does not correlate well with the physical severity of the relevant bodily disturbance.[24]

It seems that at most one of the experiences elicited by a given stimulus could be veridical/accurate under BDD, and it is unclear how we could non-arbitrarily decide which one is. Is the experience elicited the first time you are prodded with a thermal stimulator the experience that accurately tracks the level of bodily disturbance, or does the same skin site need to be prodded repeatedly before your pain system can detect the true magnitude of the disturbance? Why is the adaptation pattern different for lower vs. higher temperatures and for stimulation of the same skin site vs.

[23] In the absence of a reliable biomarker, subjective pain scales (Visual Analogue Scale, Numerical Rating Scale, Verbal Rating Scale, etc.) are commonly used to measure pain intensity in both clinical and research settings. There are obvious challenges with the use of these scales, including confounding variables such as experimenter characteristics (e.g., Kállai et al. 2004) and different understandings of verbal descriptions, e.g., "the worst pain imaginable" (Breivik 2017). Still, comparative analyses suggest that the common single-dimension scales are reliable and valid measures of felt pain intensity (Jensen 2003; Herr et al. 2004; Ferreira-Valente et al. 2011), and the replication of studies increases the validity of individual findings.

[24] It seems plausible that, even after habituation, a *higher* temperature would be correlated with *more* pain. This does little to help BDD, however. If pain experiences track mind-independent bodily disturbances and their physical features, then the same type of disturbance shouldn't be routinely associated with different kinds of pain experiences.

stimulation across different skin sites? If your pain system is a bodily disturbance detector, then what could be the reason for all this variation in pain experience; that is, what mind-independent magnitude of bodily disturbances could possibly correlate with the observed patterns?

One might argue that temporal adaptation amounts to malfunction and that the first experience elicited by a noxious stimulus is the one that accurately tracks the disturbance, but since the adaptation effects are ubiquitous and systematic and seem to occur in perfectly ordinary conditions, this response is not promising. Jepma et al. (2014) suggest that our pain sensitivity to subsequent stimuli gets dynamically updated *every time* we receive a noxious (thermal) stimulus, and that this is due to multiple adaptive processes both in the peripheral and central nervous systems. Temporal adaptation also seems beneficial to organisms. As Paul et al. write, "adapting to persistent and non-avoidable pain is an important mechanism, as it allows preserving physical, emotional, and cognitive resources" (2021, p. 1).

The proponents of BDD have made ingenious attempts at accommodating variation in experienced pain intensity. For example, Cutter and Tye (2011) argue that there is, in fact, a property of bodily disturbances that correlates with felt intensity and explains response expansion. This property is the property of "being apt to harm," indexed to a subject. Though not a physical magnitude of the bodily disturbance, Cutter and Tye suggest that it is nevertheless a natural property.[25] Supposedly this means that the property is mind-independent, even if subject-indexed, and a property that our pain experiences track (see Cutter 2017, p. 34).

It is unclear whether Cutter and Tye succeed in accommodating response expansion. But even if they did, there does not seem to be a strong enough correlation between the property of *being apt to harm* (understood as a mind-independent natural property) and experienced pain intensity to account for the patterns of temporal adaptation. For example, why would the bodily disturbance associated with repeated thermal stimulation across different skin sites be more apt to harm than the bodily disturbance associated with repeated stimulation of the same site? Cutter and Tye assume that the relevant sort of potential harm is harm "in design conditions" (2011, p. 102), and while we might be able to tell a plausible-sounding story about the harm of repeated thermal stimulation across different skin sites in our evolutionary

[25] Cutter and Tye (2011) appeal to the property of *being apt to harm* to explain both the felt intensity and the felt badness of pain. Under their view, the representational content of, *say*, the pain of stepping on a nail would be that *"*there exists a throbbing pain in my foot that is bad for me to a degree x." The property of *being bad for you* is the property of *being apt to harm you* (ibid., p. 99) and the question of whether a specific bodily disturbance is apt to harm you "depends also on facts about the organism as a whole together with facts about the environment in which the organism is embedded" (ibid., p. 101). Pautz (2014, p. 280) reads Cutter and Tye as suggesting that the aptness to harm only accounts for the differences in the unpleasantness of pains (and not intensity), but Cutter and Tye explicitly present their view as an account of variation in pain intensity (2011, pp. 104-105; see also Cutter, 2017).

environment (*i.e.,* a story about why repeated stimulation of different sites produces sensitization), this would be little more than an *ad hoc* attempt at redeeming BDD in the face of adverse evidence.[26] Gerrymandered "natural" properties do not seem very natural.

The problem gets amplified if we consider the fact that different individuals show different patterns of habituation and sensitization (Jepma et al. 2014, p. 743). Cutter and Tye's appeal to subject-indexed properties does not help, because, once again, the property of *being apt to harm* is supposed to be a "natural" property that our pain experiences track in optimal conditions. If the optimal conditions are design conditions and if the design conditions are the same for all members of the same species, then there would be *one* pattern of habituation and sensitization that optimally tracks the aptness to harm in those conditions. But the wide variation in the adaptation patterns gives little reason to think such an optimal pattern exists.

4.4 Contextual modulation of pain experience

Contextual modulation of pain experience might be the biggest problem for BDD. There is plenty of evidence that expectations, emotions, etc. can affect the experienced intensity (and perhaps also the experienced quality and location) of pain. For example, research on placebo analgesia indicates that verbal suggestions can dampen the intensity of pain (Medoff and Colloca 2015). If I offer you a sugar pill and convince you that the pill is an effective painkiller, your pain level might decrease. There are also ways to increase the felt intensity of pain. For example, showing a red light to subjects before the administration of a noxious stimulus appears to lead to the stimulus being rated more painful (Wiercioch-Kuzianik and Bąbel 2019; Moseley and Arntz 2007). If the cognitive and affective states associated with verbal suggestions (e.g., the expectation that a particular intervention will be effective) and red lights (e.g., fear or stress) routinely exert influence on the sensory dimension of pain, this presses hard on BDD. This is because, under BDD, contextual modulation presents as *interference* with the proper functioning of the pain system. Although biological systems do sometimes malfunction, it would seem odd to have a biological system so susceptible to malfunction. The specific type of malfunction also raises questions. Perceptual

[26] One could argue, for example, that noxious stimulation of different skin sites correlated with there being multiple threats (and therefore more potential for harm) in the organism's evolutionary environment, whereas repeated stimulation of the same skin site correlated with there being just one threat. That said, this sort of harm doesn't seem like a property of the bodily disturbance and therefore not something that the pain system should be able to track under BDD. Note also that we could tell a plausible-sounding story if the pattern of adaptation was reversed, *i.e.,* if repeated stimulation of the same skin site led to sensitization and stimulation of different skin sites to habituation. For example, it could be argued that repeated stimulation of the same skin site is more apt to harm e.g., by eventually burning the skin (something that wouldn't happen if different skin sites each receive just a bit of heat).

systems in general seem somewhat immune to the effects of cognitive states; some philosophers and scientists even argue that their operations are *impenetrable* to cognition (Fodor 1983; Pylyshyn 1999; see also Firestone and Scholl 2016). If pain experience is perceptual and routinely modulated by cognitive states, then we are dealing with a perceptual system that is much more vulnerable to cognitive penetration than other perceptual systems.[27] This demands an explanation.

At this point, a word of caution is in order: empirical investigation of pain experience is extremely difficult, and it is often unclear whether studies measure the effect of cognitive states on the sensory dimension of pain experience or the effects of cognitive states on other cognitive states, such as *beliefs about* sensory experience (Jacobson 2017, p. 270). Cognitive states could also influence the affective dimension of pain (*i.e.,* its unpleasantness) without influencing its perceived intensity. Nevertheless, it seems quite probable that at least *sometimes* contextual factors do influence the sensory dimension of pain. Consider this oft-cited anecdote from a medical journal:

> A builder aged 29 came to the accident and emergency department having jumped down on to a 15 cm nail. As the smallest movement of the nail was painful he was sedated with fentanyl and midazolam (…). When his boot was removed a miraculous cure appeared to have taken place. Despite entering proximal to the steel toecap the nail had penetrated between the toes: the foot was entirely uninjured. (Fisher et al. 1995, p. 70)

In the case of the builder, there is no bodily disturbance, yet the subject experiences severe pain.[28] What is the cause of the pain? It seems that the pain is caused by the person's belief that his foot is severely injured (perhaps in conjunction with distress or fear). If a belief can cause a pain experience in the *absence* of the relevant kind of bodily disturbance, then it is probable that cognitive states can also modulate pain experiences in the *presence* of noxious stimulation. Evidence from studies employing hypnotic suggestion supports this idea. For example, Hofbauer et al. (2001) used hypnotic suggestions to successfully modulate the experienced intensity

[27] The purported cases of cognitive penetration of pain experience would pose a problem for BDD even if some of them were explainable in terms of *sensory* inputs to the pain system. For example, Casser and Clarke suggest that "…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" (2022, p. 15). If fear causes the organism's heart rate and/or stress level to increase, it could modulate pain experience *indirectly* (which might not count as cognitive penetration proper). But unless the heuristic processes aim at deciphering levels of bodily disturbance (this is not what Casser and Clarke argue), they still present as interference with the normal and proper functioning of the pain system under BDD. Furthermore, it is not at all obvious that any/all purported cases of cognitive penetration of pain experience are explainable in terms of such sensory inputs alone.

[28] Another example might be the phenomenon of *vicarious* pain where one feels pain in their own body when witnessing (or imagining) others undergo painful experiences (Fitzgibbon et al. 2010).

of a painful heat stimulus, and the resulting changes in pain intensity corresponded to changes in cortical activation patterns.[29]

Tye (1995a) acknowledges that relaxation and enjoyment can dampen pain experience and that anxiety can heighten it but argues that such effects do not challenge his view:

> These facts…may be explained by supposing that the brain receptor pathway in the spinal column leading to the somatosensory cortex (…) has a gate in it that is controlled from the higher brain centers (the gate control theory). When this gate is partly closed, less information gets through, and the feeling of pain diminishes. As it opens further, more information is enabled to pass. Anxiety, excitement, joy, concentration, and other higher-level activities affect the orientation of the gate. So, the fact that the experience of pain is, *in the above sense*, cognitively penetrable presents no real difficulty for my proposal. What happens is simply that one's cognitive assessment of the situation feeds back down into the sensory module for the experience of pain and *affects how much information gets through about bodily damage.* (1995a, pp. 114-115; second italics mine)

Tye's response is puzzling. In cases of top-down modulation of pain experience, the bodily disturbance stays the same but the experience changes. It is true that this variation likely results from the excitatory and inhibitory control of nociceptive channels, but it is not at all clear how this is consistent with BDD. To start, Tye mischaracterizes the issue. The issue is not the *amount* of information about the bodily disturbance that is let through, but the *accuracy* of the resulting experience. If you are feeling anxious and a mild disturbance gives rise to an excruciatingly intense pain experience, then the experience *mis*represents the qualities of the disturbance, *according to Tye's own view*. Recall Tye's contention that a "twinge of pain represents a mild, brief disturbance" (1995a, p. 113). Could a sudden excruciating stab of pain represent the same mild bodily disturbance? It is difficult to see how this could be if the pain system is in the business of tracking mind-independent bodily disturbances and their features. It is even more difficult to see how an intense pain would provide more information about a mild disturbance than a mere twinge of pain.

[29] This study by Hofbauer et al. (2001) and another study by Rainville et al. (1997) together support the view that the processing of the sensory and affective dimensions of pain involve different cortical areas. When hypnotic suggestions were directed at altering experienced pain intensity, this corresponded to modulation of pain-evoked activity in the somatosensory cortices but not in the anterior cingulate cortex (Hofbauer et al. 2001). When the suggestions were directed at altering the unpleasantness of pain, significant modulation of pain-evoked activity was observed in the anterior cingulate cortex but not in the somatosensory cortices (Rainville et al. 1997).

Hill (2017) acknowledges that the experienced intensity of pain correlates better with neural states than it does with bodily disturbances (or the activity of the nociceptors which respond to those disturbances). As a result, he embraces the conclusion that pain experiences systematically *mis*represent bodily disturbances. More specifically, he maintains that pain experiences mistakenly attribute properties of neural states to bodily disturbances (ibid., p. 66). Though interesting, this response is not available to the proponent of BDD; the detecting/tracking of neural properties is not consistent with the requirement that the contents of pain states be mind-independent. But even on a more general level, it makes little sense to conceptualize the pain system as a bodily disturbance detector if pain experiences systematically misrepresent those disturbances. Consider a simple Geiger counter with an analog readout. If the measurements indicated on the readout weren't correlated with the levels of ionizing radiation in the environment (but were correlated with some internal states of the instrument instead), the only reason to call the Geiger counter a radiation detector would be the fact that it was *designed* (badly!) for that purpose. But with the pain system we cannot simply *assume* that it was "designed" to be a bodily disturbance detector; we need some reason to think that this is the case. If the relationship between bodily disturbances and pain experiences is highly variable and if there is evidence of substantial contextual modulation of pain experience, such a reason is lacking.

For Hill (2017) it suffices if pains are bodily disturbances that our pain experiences represent, even if those experiences systematically *mis*represent some features of the disturbances. That said, if the pain system is not in the business of detecting/tracking bodily disturbances, it is not clear why we should identify pains with those disturbances in the first place. What motivates this view of the nature of pain, if we are forced to give up on the idea that there is a clear correlation with the experiential state and the bodily state it supposedly represents?

So far, I have sought to show that the detector view of pain falls short when it comes to the task of accommodating and explaining systematic pain phenomena. Since the ability to accommodate and explain relevant phenomena is a key desideratum for (empirically informed) philosophical theories, I conclude that the detector view should be abandoned in favor of a better-fitting conception of the function of pain. In the next section, I propose such a conception.

5 The pain system as a sophisticated security system

If the pain system is not a bodily disturbance detector, then what is it? A comprehensive answer is beyond the scope of the present work, but I hope to lay the groundwork for an answer here. The first step is to replace the old detector metaphor

with a new, better-fitting one. The metaphor I have in mind is that of a *sophisticated security system*. This metaphor is not just a metaphor, of course; it is an alternative conception of the function of the pain system, one that better explains and accommodates the empirical data considered above.

5.1 The pain system and detectors

Recall the two striking examples of the variable relationship between bodily disturbances and pain experiences: painless injury and pain experience in the absence of bodily disturbances. BDD has trouble accounting for these phenomena because it maintains that the pain system aims to track/detect disturbed states of our body. Tracking/detection of bodily disturbances requires correspondence between the experience and the relevant bodily state, and such correspondence is clearly lacking in these cases.

The pain system *involves* detectors, most notably peripheral nociceptors. The function of the nociceptors might even be to detect bodily disturbances. If the activation of nociceptors correlated strongly and straightforwardly with pain experience, this would afford a great deal of credibility to BDD.[30] Unfortunately, however, such correlation is lacking, and some neuroscientists have even suggested that nociception and pain perception are two distinct processes that should be studied separately (Apkarian 2017, p. 74).[31] Though nociceptors are generally taken to be the "sensors of the pain pathway," the pathways themselves are modulated by both peripheral and central mechanisms (Dubin and Patapoutian 2010). That this modulation appears to be part of the normal functioning of the pain system strongly suggests that the pain system *as a whole* is not a detector.

In addition, the pain system appears to receive inputs not just from nociceptors but also from vision, audition, interoception (heart rate, stress level, etc.), cognition, etc. I mentioned earlier that red light might increase the felt intensity of noxious stimulation. When the pain system is conceptualized as a bodily disturbance detector, such effects seem quite mysterious. After all, red lights in themselves do not carry information about bodily disturbances or make bodily disturbances worse. But if the pain system is conceptualized as a sophisticated security system, it makes

[30] The intuitive idea that nociception and pain experience are strongly correlated goes back to Descartes. In *Treatise on Man*, Descartes conjectures that there exists a gate between the brain and the sensory nerves that can be opened by the "pulling" of the fibers in those nerves by external objects. In the case of pain, somatosensory fibers are supposedly pulled so forcibly that they become severed. When the gate in the brain is opened, this "gives occasion for the soul…to have the sensation of *pain*" ([1664] 1985), p. 103). Though Descartes differentiates between the mechanical transmission of the sensory signals and the perceptual experience of pain, he also thinks that God has ensured that "the soul will have different sensations corresponding to the different means in which the entrances…are opened by means of nerves" (ibid., p. 102).

[31] But see Pereplyotchik 2017.

sense for it to be responsive to information about threats to the organism. Since red color is often associated with danger, it could be taken to indicate threat.

5.2 The pain system and dynamic updates

Detectors do not dynamically update their sensitivity levels, or at least they are not supposed to. Consider a simple thermometer. The job of a thermometer is to accurately measure temperature and in optimal conditions this is what a decent thermometer does. Placing an ambient air thermometer in direct sunlight might cause it to measure a temperature higher than that of the surrounding air, but these are *not* the conditions in which the thermometer can be expected to function properly. In optimal conditions, the same air temperature should always give rise to the same reading.

Now consider a simple home security system that has customizable sensitivity settings for, *say*, motion detectors at windows. If you live with cats, you might not want an alarm to go off if they jump on the window. But if the house is empty, lowering the settings might be a good idea, since motion near your windows is now more likely to indicate danger. I suggest that the pain system bears more affinities to a home security system than it does to a thermometer. There are no optimal conditions in which the same noxious stimulation or the same bodily disturbance would reliably give rise to a specific type of pain experience. This is because the pain system, like the home security system, has "customizable" sensitivity settings. These settings are updated in real time through temporal adaptation and contextual modulation, making the pain system a very *sophisticated* kind of security system, responsive to different types of evidence of danger.

Stress, anxiety, and fear can provide evidence of danger. So can visual information. Studies employing virtual reality make this apparent. One such study found that pain experience can be evoked by the mere *appearance* of being touched (Harvie et al. 2022). When a subject suffering from complex regional pain syndrome (CRPS) put his hand under a virtual tap simulating running water, this "evoked instant and severe 'sharp' pain" and "an immediate withdrawal response" (ibid., p. 1875). One of the symptoms of CRPS is allodynia, and even a light touching of the affected limb can be painful. For a person suffering from CRPS, any contact with the affected limb might therefore come to mean danger, and visual evidence of such contact might be sufficient to produce a pain response.

Although CRPS is an abnormal condition that likely involves changes to peripheral and/or central nervous systems (usually following a forceful trauma to the affected limb), the case above is in many ways analogous to the case of the 29-year-old

builder.[32] Both scenarios involve pain experiences that seem ultimately "useless"—virtual reality does not pose a real threat to one's body and neither does a nail piercing one's shoe if it misses one's foot—but they can provide clues as to how the pain system normally operates.[33]

5.3 The pain system and action guidance

Detectors detect, security systems protect. Particularly smart home security systems might, for example, connect to locks, doors, thermostats, and appliances, and regulate their operations with the goal of protecting the building and its inhabitants. The outputs of the pain system are obviously much more fine-grained than those of any home security system on the market. They can also guide your behavior in manifold ways. If you burn your finger, you feel a type of pain that causes you to quickly withdraw your hand. If you have a migraine, your pain might prompt you to limit your exposure to visual, auditory and/or somatosensory stimuli. If you sprain your ankle, your pain tells you not to put weight on the ankle. In all these cases, the pain itself appears to be guiding your behavior. And not only that; the pain also helps you *learn* to behave in a way that protects your bodily integrity in the future (a child who touches a hot stove quickly learns not to do it again!). On the other hand, a stimulus that used to be painful might cease to give rise to pain experiences if we learn that the stimulus is ultimately harmless. This is the guiding idea behind many chronic pain interventions, such as pain neuroscience education (see, e.g., Louw et al. 2016).[34]

One way to flesh out the idea that pain is directly action-guiding is to say that the function of pain is to enable and enhance the manifestation of certain kinds of competences. These competences are tied to the survival and health of the organism. For example, the pain of a sprained ankle might help a person develop and manifest a recuperative (healing-promoting) competence, and the pain of touching a hot stove might help a person develop and manifest a danger avoidance competence.

[32] Other studies support this idea. For example, Acerra and Moseley (2005) used a mirror procedure to create a virtual representative of a limb affected by CRPS. The unaffected limb of the participant was then stimulated while the participant focused on its mirror image (*i.e.,* the "virtual stand-in" for the affected limb). As the authors describe, "when the stimulation site on the pain-free limb corresponded to an area of allodynia on the affected limb, CRPS1 patients reported normal sensation at the stimulated site but pain at the corresponding site on the affected limb. Several patients quickly withdrew their affected limb, and two patients chose not to finish an assessment because it was too painful" (ibid., p. 751).

[33] It might be that in these cases the visual information produces anxiety or fear, making the emotional state the direct cause of the pain experience. But it could also be that vision affects pain processing more directly. For example, Harvie et al. write that "The fact that pain was visually evoked implies that bimodal visual-tactile and/or mirror neurones may be involved and solely activated by vision" (2022, p. 1877).

[34] This makes pain perception different from paradigmatic perceptual modalities. Though there are critical periods in the development of perceptual systems, it does not seem likely that our responses to e.g., visual stimuli could be radically altered through learning later in life.

Recuperative competence is manifested when an injury is allowed to heal quickly and efficiently. This might involve resting the injured body part or giving it the care that it needs. Danger avoidance competence, on the other hand, is manifested when contact with a potentially threatening or damaging environmental object/event is cut short or altogether avoided.

Recuperative competence and danger avoidance competence are *behavioral* competences. But it seems that pain can also help us develop and manifest *cognitive* competences. One example is danger recognition. Note that pain experiences seem to guide future decision-making and risk analysis. When pain experiences are absent, people struggle with these tasks. Congenital insensitivity to pain (CIP) is a case in point. Individuals with CIP either never develop nociceptors or their nociceptors fail to respond to noxious stimulation. As a result, the affected individuals never experience physical pain (Drissi et al. 2020, p. 67). Because of repeated injuries, self-mutilation, and poor healing outcomes, significant morbidity—and often untimely death—follows. Though people with CIP might correctly *believe* that they should regularly check for injuries, soft-round sharp edges in their homes, allow their injuries to properly heal, and never put their hand in boiling water, without pain experience, such beliefs are often ineffectual. What CIP shows is that pain is critical when it comes to learning to moderate behavior and analyze risks.

5.4 The sophisticated security system view of pain

Putting together these different strands (modulation of nociceptive inputs, dynamic updating of sensitivity levels, action guidance), we arrive at a promising preliminary characterization of the sophisticated security system view of pain:

> [SSS] The pain system is a sophisticated security system: the function of pain states is to enhance and enable the manifestation of behavioral competences (e.g., danger avoidance) and cognitive competences (e.g., danger recognition) in a context-responsive manner.

If SSS is on the right track, then pain experiences have success conditions; a pain experience is successful if it helps an organism manifest some relevant competences and unsuccessful if it doesn't. This means that the pain system can malfunction. One example of such malfunction is chronic pain that presents without protective or recuperative benefits. The pain is maladaptive because it often leads the affected individual to avoid behaviors that are health-promoting and to engage in behaviors that increase morbidity. That said, because the pain system is not a detector of mind-independent states of the body, variation in pain intensity or pain thresholds need not indicate malfunction. The same bodily disturbance can be more dangerous in

some circumstances than it is in others. If I lack access to social support, I am sleep-deprived, or my situation is otherwise precarious, a twisted ankle might be dangerous. If, on the other hand, support is available, I am well-rested, and my situation is generally safe, it might not be.

6 Further developments

In this section, I discuss how SSS might guide further philosophical explorations of pain, how a proponent of the conception might respond to some worries that immediately arise, and how the conception might have practical relevance in clinical and research settings.

6.1 Philosophy of pain: the function question, the content question, and the nature question

Many contemporary philosophers of pain are concerned with the function of pain (the function question), the contents of pain experience (the content question),[35] and the nature of pain (the nature question). Many of these philosophers are also committed to an empirically informed approach to answering these questions. I submit that the function question is a particularly useful starting point for empirically informed philosophical explorations of pain. This is because the evidence from pain science, especially psychophysical evidence, has more *direct* relevance to answering the function question than it does to answering the content and nature questions; data on how the pain system works can ground inferences about how the system is "designed" to work.

SSS acknowledges that the subjects of pain are organisms embedded in meaningful environments with experiences and histories that color their interactions with those environments. As a result, the view can easily accommodate and explain the pain phenomena discussed in this chapter. Painless injuries occur because sometimes pain might hinder the manifestation of relevant competences. In such cases, the system inhibits the transmission of nociceptive signals and doing so is perfectly consistent with the function of the system, which is to protect the organism.[36] Pain experience in

[35] I understand the phrase "the contents of experience" broadly without a commitment to accuracy conditions.

[36] A relevant difference between the SSS explanation of adaptive, SIA-mediated painless injury and the BDD kill switch explanation of the same phenomenon (see Sect. 4.1) is that the latter portrays painless injury as a case where the pain system is *not* doing its job, whereas the former can maintain that the pain system is functioning properly even if it is suppressing pain experience. The proponent of BDD could argue that there is nothing wrong with the idea of the system working against its goal in some special circumstances (where it is beneficial for the organism if it does so) but recall that painless injury is but an extreme example of the kind of modulation that

the absence of a bodily disturbance is evidence of system malfunction, but it also suggests that the system is highly plastic to begin with. This plasticity also manifests as temporal adaptation and contextual modulation of pain even in perfectly healthy individuals in perfectly ordinary situations. The plasticity lends support to the idea that the function of pain is to keep the organism safe in *changing* conditions.

Ultimately, answers to the content and nature questions are also owed but by answering the function question first, we can effectively limit the number of available answers to these very difficult problems. For example, my conceptualization of the pain system as a sophisticated security system suggests that pains do not have descriptive or indicative contents of the type "there exists a disturbance (with such-as-such features) in this bodily location." This leaves the door open for other kinds of descriptive contents (e.g., "there is (this much) threat to my body in this particular situation")[37] and/or imperative contents (e.g., "protect your finger!").[38] This means that SSS can be made consistent with the view that the phenomenal character of pain supervenes on its content. Crucially, SSS does *not* entail that pain experience isn't body-directed.

Relatedly, SSS suggests that pain should not be analyzed as a disturbed state of the body or as an experience perceptually representing such a state.[39] But, again, this leaves the door open for many other kinds of answers to the nature question. Pain could be a *command* issued by the body to protect a body part (Klein 2015), a type of *emotion* (Craig 2003), or a felt, context-dependent *estimate of bodily threat* (see Kiverstein et al. 2022). These options are rooted in the notion of pain as a *homeostatic sensation* motivating situationally appropriate actions to eliminate threats to one's bodily integrity (see Klein 2015, pp.13-16; Craig 2003, p. 304). As such, they are highly compatible with the sophisticated security system view of pain. That said, other answers to the nature question could also harmonize with SSS. For example, pain could be conceptualized as a complex biological process with multiple levels or dimensions, including experiential ones (see, e.g., Smith et al. 2011; Hardcastle 1999).

---

nociceptive signals undergo in perfectly ordinary contexts. This modulation either is consistent with the function of the pain system or isn't. SSS says it is, BDD cannot.

[37] Kiverstein et al. advocate an "embodied predictive processing" theory of pain and suggest that the function of pain is "to move the person to take urgent action to avoid further, possibly life-threatening damage to the body" (2022, p. 979). According to the authors, "pain is the outcome of unconscious inferential processes that are distributed across the homeostatic processes that make up the person's body" (ibid., p. 974). But rather than reflect estimates of bodily disturbances, pain experiences reflect "the 'overall estimate' of threat that is posed to the body in a particular environment…" (ibid., p. 987).

[38] In other words, the view is consistent with some forms of pain imperativism, especially Klein's (2015). For a version of imperativism that has affinities with BDD, see Martínez 2011.

[39] This idea is not new. Consider, for example, the words of Patrick Wall: "Pain is taken not as a simple sensory experience signalling the existence of damaged tissue. The presence and intensity of pain is too poorly related to the degree of damage to be considered such a messenger." (1979, p. 264) With decades of additional research, the case against both the bodily disturbance view and the messenger view can now be made even stronger.

6.2 Pain, beliefs, and perception

Pain experiences appear to play a role in the formation of beliefs about bodily disturbances. If I feel a sharp, burning, well-localized pain in my arm, I might correctly infer that I have been stung by a bee and that its venom is now entering my tissues. This could be taken to imply that, contrary to what I have argued here, pain experiences *do* function to inform us about bodily disturbances.[40] So the question is: can SSS accommodate the role that pain experience plays in the formation of beliefs about the disturbed states of our bodies?

It is true that pains can often be used to infer that *something* is disturbed in our bodies. But when it comes to more fine-grained information, pain doesn't seem to tell us much. Klein (2015, p. 2) suggests that pains are, in fact, strikingly uninformative, and I am inclined to agree. For example, If I feel a sudden pain in my leg, I normally need to look to know what is going on (or perhaps feel the leg with my hand). It is difficult to tell what causes the pain or how bad the injury is *based on the pain experience alone*. I suspect that it is easy to forget just how uninformative pain really is because pain prompts us to immediately search for more information about its cause. A pain experience tells us that something is wrong or that something needs to be done, but after it has accomplished this, we tend to rely on visual, tactile, proprioceptive, etc. information to fill in the gaps.

It might even be that the spatial content of pain experience is rather uninformative when there is no help available from proprioception, touch, or other relevant sensory systems. For example, Skrzypulec (2021) argues that we have good reason to think that pain experiences do not involve field-like content similar to that of tactile experiences. He cites studies by Mancini et al. (2015), in which two different skin sites of participants were simultaneously stimulated by lasers (amounting to noxious stimulation in the absence of tactile stimulation) and in which the participants were asked to estimate the distance between the evoked sensations. As Skrzypulec recounts, the participants' estimates were all over the place.[41]

Skrzypulec does not claim that pain has no spatial content; he thinks that pain sensations involve both feature-placing content (representing that there is pain in

[40] For example, Hill writes that "if higher cognitive agencies are to protect the organism from harm, and prevent such harms from becoming worse, they need an accurate, continuously updated map of the peripheral loci where there is actual or potential damage" (2009, p. 180).

[41] Skrzypulec writes: "In some participants, greater physical distance corresponded to greater perceived distances; in some participants, perceived distances were virtually constant despite changes in physical distances; and in some participants, the correlation was negative such that greater physical distance corresponded to *smaller* perceived distances" (2021, p. 563).

some location of the body) and identity content (representing that two pains are spatially *distinct*). What he does suggest is that the spatial content of pain experiences is less developed than the content of interoceptive tactile experiences. The fact that noxious stimulation and tactile stimulation often co-occur in paradigmatic pain cases might explain our tendency to attribute more detailed spatial content to our pain experience; we aren't normally aware of the contributions of non-nociceptive sensory information to our pain-involving experiences.

Second, as Hill (2009) notes, the experiential awareness of pain has many similarities with paradigmatic forms of perception such as vision: pain experience has a proprietary phenomenology, it is particularized and localized (no matter how poorly), etc. In view of these similarities, argues Hill, "it would be very uncomfortable to withhold the label "perceptual system" from the structures that subserve awareness of pain" (2009, p. 174).[42] Assuming that Hill is right, it seems that we have good reason to be suspicious of any conception of pain which entails that the pain system isn't a perceptual system. So, does SSS entail that?

The answer to this question depends largely on what we take to be the goal of perception. If perception is supposed to be in the business of representing the mind-independent world as it is, then pain experience is not perceptual under SSS. But it is not clear that paradigmatic perceptual experiences are any more "perceptual" *in this sense*. For example, many philosophers have argued that the function of color vision is not to detect physical properties but to present the world in a way that guides successful action (Hatfield 1992; Thompson 1995; Chirimuuta 2015; Rosenqvist 2023b). And, on a more general level, the conception of perception as tracking and representing mind-independent properties of the environment has been

[42] As mentioned, Hill (2009) is a proponent of a more complex version of the bodily disturbance detector view. He states that pain experiences have the "function of encoding information" about bodily disturbances (2009, p. 180) but maintains that such experiences do not accurately track *all* aspects of the disturbances, e.g., intensity. As such, the view seems consistent with the possibility that pain systems have *additional* functions, perhaps of the sort that I have proposed in the present work, even if Hill himself does not advocate for this possibility. Nevertheless, such a hybrid view would (i) still entail that the pain system is a very *bad* bodily disturbance detector and (ii) because the correlation between the severity of bodily disturbances and the intensity of pain experiences is so weak, we would still need other reasons for thinking that the pain system has the function of detecting such disturbances in the first place. Hill himself finds a reason in the above-mentioned similarities between pain experiences and paradigmatic perceptual experiences; a key motivation for his view is the observation that pain experiences represent pains as having specific bodily locations and extents (pp. 177, 179). I have discussed evidence suggesting that the spatial content of pain does not include specific locations or extents. But even if pain experiences *did* assign specific locations and extents to pains, it is nevertheless not true, as Hill suggests, that apart from bodily disturbance detector views, "every other theory implies that our experiences of pain are systematically misleading, or illusory, insofar as they assign locations to pains" (p. 178). For example, Klein's (2015) imperativist theory of pain takes it at face value that pains have locations, without identifying pains with bodily disturbances or maintaining that pain systems have the function of detecting and encoding information about such disturbances. Finally, I have argued that the sophisticated security system view can account for all the relevant pain phenomena, including the paradigmatic cases, which suggests that a hybrid view would be unmotivated and unnecessarily complex.

challenged by scholars across disciplines (Hatfield 2009; Hoffman et al. 2015; Koenderink 2014; Purves et al. 2015; Felin et al. 2017). If the goal of perception is simply to guide successful action, then there is no inconsistency. Pain experience is body-directed and action-guiding, just like color experience is external-directed and action-guiding.[43]

On the other hand, there are also relevant differences between pain experience and the paradigmatic forms of perceptual experience. For example, pain experience seems more directly motivating than visual experience. Klein suggests that this is because pains have imperative (and *not* indicative) contents, whereas visual experiences have indicative (and *not* imperative) contents (2015, p. 3 and Chs. 5-6). But we could also explain this difference in terms of the competences that the relevant perceptual systems help organisms manifest. For example, (human) color visual systems might be primarily involved in the manifestation of *perceptual* competences (Rosenqvist 2023a), whereas pain systems might help organisms manifest *behavioral* and *cognitive* competences more directly. If this were the function of pain, then it would (partially) explain the stronger motivational force.

6.2 Practical applications

The function question also has practical relevance in clinical and research contexts. Consider, for example, the current International Association for the Study of Pain (IASP) definition of pain as "[a]n unpleasant sensory and emotional experience associated with, or resembling that associated with, *actual or potential tissue damage*" (Raja et al. 2020, p. 1977; italics mine). The IASP brings together "scientists, clinicians, health care providers, and policymakers from around the world *in pursuit of their mission to bring relief to those who are in pain.*"[44] However, the emphasis of the IASP definition on actual or potential tissue damage might have unintended negative consequences that come in the way of accomplishing this goal. For example, Smith et al. (2011) trace some of the "barriers to improved [pain] classification and subsequent research" to traditional perspectives and assumptions, such as those implicit in the IASP definition. The authors advocate amendments to the IASP definition (such as dropping the reference to tissue damage) to facilitate research into, and treatment of, chronic pain disorders which manifest with pain disproportionate to the observed bodily disturbances.[45] By drawing attention to the

[43] This is not to say that pain could not also be external-directed in some sense. For example, Kiverstein et al. propose that pain is simultaneously "a local sensing of the body and… a more global, all-encompassing attitude towards the environment" (2022, p. 975).
[44] See the official website of the IASP: https://www.iasp-pain.org/about/ (accessed Oct 24, 2023; italics mine).
[45] Smith et al. (2011) work with the previous IASP definition, according to which pain is "an unpleasant sensory and emotional experience associated with actual or potential tissue damage, or described in terms of such damage" (Merskey et al. 1979), but some of their criticisms also apply to the revised version, which retains the emphasis on tissue damage.

complexity of the pain system and by conceptualizing the function of the system without explicit reference to bodily disturbances, SSS could have similar benefits.

When it comes to clinical settings, Johnson et al. (2023) connect the use of pain metaphors aligning with "explanatory models associating pain with actual or potential tissue damage" to increased persistence of pain, and advocate a shift towards "constructive, holistic, and person-centred" pain language to promote better patient outcomes. Not only does SSS explain how the choice of metaphor might influence the intensity and duration of pain (the use of negative metaphors could be seen as providing evidence of danger) but it also directly contributes to the development of more constructive and person-centered pain language.

7 Conclusion

If you step on a nail, you will likely experience pain. But when it comes to the specific features of that pain, it is difficult to predict what you will experience based on the physical condition of your foot alone. This is because the relationship between pains and bodily disturbances is messy. This messy relationship, together with the observation that the pain system is responsive to various kinds of evidence of danger and safety, suggests that the pain system is not a bodily disturbance detector, but a sophisticated security system. Instead of aiming to inform us about the mind-independent disturbances to parts of our body, our pain systems simply aim to keep us safe.

I hope that my brief sketch of the sophisticated security system view is enough to show its potential when it comes to facilitating our growing understanding of pain. Empirical data from the pain sciences is crucial, but providing a coherent philosophical account of pain can also have far-reaching consequences. Pain is a public health challenge and a leading cause of disability worldwide. It is difficult to study, difficult to communicate, and difficult to treat. At least some of these challenges have to do with how the pain system is conceptualized.

References

Acerra NE, Moseley GL (2005) Dysynchiria: watching the mirror image of the unaffected limb elicits pain on the affected side. Neurology 65:751–753. 10.1212/01.wnl.0000178745.11996.8c

Apkarian V (2017) Advances in the neuroscience of pain. In: Corns J (ed) The Routledge Handbook of Philosophy of Pain. Routledge, London, p 73–86

Aydede M (2009) Is feeling pain the perception of something? Journal of Philosophy 106:531–567. https://doi.org/10.5840/jphil20091061033

Aydede M (2017) Pain: Perception or introspection? In: Corns J (ed) The Routledge Handbook of Philosophy of Pain. Routledge, London, p 221–232

Bain D (2013) What Makes Pains Unpleasant? Philosophical Studies (Supplement) 166:69–89. https://doi.org/10.1007/s11098-012-0049-7

Bartley EJ, Fillingim RB (2013) Sex differences in pain: a brief review of clinical and experimental findings. British Journal of Anaesthesia 111:52–58. https://doi.org/10.1093/bja/aet127

Beecher HK (1956) Relationship of significance of wound to pain experienced. Journal of the American Medical Association CLXI:1609–1613. 10.1001/jama.1956.02970170005002

Breivik H (2017) Patients' subjective acute pain rating scales (VAS, NRS) are fine; more elaborate evaluations needed for chronic pain, especially in the elderly and demented patients. Scandinavian Journal of Pain 15:73–74. https://doi.org/10.1016/j.sjpain.2017.01.001

Burge T (2010) Origins of Objectivity. Oxford University Press, Oxford

Butler RK, Finn DP (2019) Stress-induced analgesia. Progress in Neurobiology 88:184–202. https://doi.org/10.1016/j.pneurobio.2009.04.003

Byrne A (2001) Intentionalism Defended. Philosophical Review 110:199–240. https://doi.org/10.2307/2693675

Byrne A, Hilbert D (2003) Color realism and color science. Behavioral and Brain Sciences 26:3–21. https://doi.org/10.1017/S0140525X03000013

Casser LC (2021) The Function of Pain. Australasian Journal of Philosophy 99:364–378. https://doi.org/10.1080/00048402.2020.1735459

Casser L, Clarke S (2022) Is pain Modular? Mind & Language:1–19. https://doi.org/10.1111/mila.12430

Chirimuuta M (2015) Outside Color: Perceptual Science and the Puzzle of Color in Philosophy. MIT Press, Cambridge, MA

Cohen J (2009) The Red & The Real: An Essay on Color Ontology. Oxford University Press, Oxford

Craig AD (2003) A new view of pain as a homeostatic emotion. Trends in Neurosciences 26(6):303–307. https://doi.org/10.1016/S0166-2236(03)00123-1

Cummins R (1975) Functional Analysis. The Journal of Philosophy 72:741–765. https://doi.org/10.2307/2024640

Cutter B, Tye M (2011) Tracking Representationalism and the Painfulness of Pain. Philosophical Issues 21:90–109. https://doi.org/10.1111/j.1533-6077.2011.00199.x

Cutter B (2017). Pain and Representation. In: Corns J (ed) The Routledge Handbook of Philosophy of Pain. Routledge, London, p 29–39

deCharms RC, Maeda F, Glover GH, Ludlow D, Pauly JM, Soneji D, Gabrieli JDE, Mackey SC (2005). Control over brain activation and pain learned by using real-time functional MRI. Proceedings of the National Academy of the Sciences U.S.A. 102:18626–18631. 10.1073/pnas.0505210102

Descartes R ([1664] 1985) Treatise on Man. In: The Philosophical Writings of Descartes Volume 1 (trans: Cottingham J, Stoolholft R, Murdoch D). Cambridge University Press, Cambridge, p 99–108

Dretske F ([1995] 1997) Naturalizing the Mind. MIT Press, Cambridge, MA

Drissi I, Woods WA, Woods CG (2020) Understanding the genetic basis of congenital insensitivity to pain. British Medical Bulletin 135:65–78. 10.1093/bmb/ldaa003

Dubin AE, Patapoutian A (2010) Nociceptors: the sensors of the pain pathway. The Journal of Clinical Investigation 120:3760–3772. 10.1172/JCI42843

Felin T, Koenderink J, Krueger JI (2017) Rationality, perception, and the all-seeing eye. Psychonomic Bulletin & Review 24:1040–1059. https://doi.org/10.3758/s13423-016-1198-z

Ferreira-Valente MA, Pais-Ribeiro JL, Jensen MP (2011) Validity of four pain intensity rating scales. Pain 152(10):2399–2404. 10.1016/j.pain.2011.07.005

Firestone C, Scholl BJ (2016) Cognition does not affect perception: Evaluating the evidence for "top-down" effects. Behavioral and Brain Sciences 39:e229. 10.1017/S0140525X15000965

Fisher JP, Hassan DT, O'Connor N (1995) Minerva. British Medical Journal 310:70. https://doi.org/10.1136/bmj.310.6971.70

Fitzgibbon BM, Giummarra MJ, Georgiou-Karistianis N, Enticott PG, Bradshaw JL (2010) Shared pain: from empathy to synaesthesia. Neuroscience & Biobehavioural Reviews 34:500–512. https://doi.org/10.1016/j.neubiorev.2009.10.007

Fodor JA (1983) The Modularity of Mind. MIT Press, Cambridge, MA

Goksan S, Hartley C, Emery F, Cockrill N, Poorun R, Moultrie F, Rogers R, Campbell J, Sanders M, Adams E, Clare S, Jenkinson M, Tracey I, Slater R (2015) fMRI reveals neural activity overlap between adult and infant pain. eLife 4:e06356. https://doi.org/10.7554/eLife.06356

Hanyu-Deutmeyer AA, Cascella M, Varacallo M (2022) Phantom Limb Pain. In: StatPearls [Internet]. StatPearls Publishing, Treasure Island. https://www.ncbi.nlm.nih.gov/books/NBK448188/

Hardcastle VG (1999) The Myth of Pain. MIT Press, Cambridge, MA

Hardin CL ([1988]1993) Color for Philosophers: Unweaving the Rainbow. Hackett Publishing Company, Indianapolis

Harvie DS, Stanton TR, Kennedy H, Coppieters MW (2022) Visually evoked pain and its extinction using virtual reality in a patient with complex regional pain syndrome type II. Pain 163:1874–1878. 10.1097/j.pain.0000000000002605

Hatfield G (1990) The Natural and the Normative: Theories of Spatial Perception from Kant to Helmholtz. MIT Press, Cambridge, MA

Hatfield G (1992) Color Perception and Neural Encoding: Does Metameric Matching Entail a Loss of Information? PSA; Proceedings of the Biennial Meeting of the Philosophy of Science Association 1:492–504. https://doi.org/10.1086/psaprocbienmeetp.1992.1.192778

Hatfield G (2009) On Perceptual Constancy. In: Hatfield G, Perception and Cognition: Essays in the Philosophy of Psychology. Clarendon Press, Oxford, p 178–211

Herr KA, Spratt K, Mobily PR, Richardson G (2004) Pain intensity assessment in older adults: Use of experimental pain to compare psychometric properties and usability of selected pain scales with younger adults. Clinical Journal of Pain 20(4):207–219. 10.1097/00002508-200407000-00002

Hill CS (2009) Consciousness. Cambridge University Press, Cambridge

Hill CS (2017) Fault lines in familiar concepts of pain. In: Corns J (ed) The Routledge Handbook of Philosophy of Pain. Routledge, London, p 60–69

Hofbauer RK, Rainville P, Duncan GH, Bushnell MC (2001) Cortical Representation of the Sensory Dimension of Pain. Journal of Neurophysiology 86:402–411. https://doi.org/10.1152/jn.2001.86.1.402

Hoffman DD, Singh M, Prakash C (2015) The interface theory of perception. Psychonomic Bulletin and Review 22:1480–1506. https://doi.org/10.3758/s13423-015-0890-8

Humes LE, Busey TA, Craig JC, Kewley-Port D (2009) The effects of age on sensory thresholds and temporal gap detection in hearing, vision, and touch. Attention, Perception, & Psychophysics 71:860–871. https://doi.org/10.3758/APP.71.4.860

Jacobson H (2017) Pain and Cognitive Penetrability. In: Corns J (ed) The Routledge Handbook of Philosophy of Pain. Routledge, London, p 266–275

Jensen MP (2003) The Validity and Reliability of Pain Measures in Adults with Cancer. The Journal of Pain 4(1):2–21. 10.1054/jpai.2003.1

Jepma M, Jones M, Wager TD (2014) The dynamics of pain: Evidence for simultaneous site-specific habituation and site-nonspecific sensitization in thermal pain. Journal of Pain 15:734–746. 10.1016/j.jpain.2014.02.010

Johnson MI, Hudson M, Cormac GR (2023) Perspectives on the insidious nature of pain metaphor: we literally need to change out metaphors. Frontiers in Pain Research 4:1224139. https://doi.org/10.3389/fpain.2023.1224139

Julien N, Goffaux P, Arsenault P, Marchand S (2005) Widespread pain in fibromyalgia is related to a deficit of endogenous pain inhibition. Pain 114:295–302. https://doi.org/10.1016/j.pain.2004.12.032

Kállai I, Barke A, Voss U (2004) The effects of experimenter characteristics on pain reports in women and men. Pain 112(1-2):142–147. 10.1016/j.pain.2004.08.008

Kiverstein J, Kirchhoff MD, Thacker M (2022) An Embodied Predictive Processing Theory of Pain Experience. Review of Philosophy and Psychology 13:973–998. https://doi.org/10.1007/s13164-022-00616-2

Klein C (2015) What the Body Commands: An Imperative Theory of Pain. MIT Press, Cambridge, MA

Koenderink J (2014) The All Seeing Eye? Perception 43:1–6. https://doi.org/10.1068/p4301ed

Lautenbacher S, Peters JH, Heesen M, Scheel J, Kunz M (2017) Age changes in pain perception: A systematic-review and meta-analysis of age effects on pain and tolerance thresholds. Neuroscience & Biobehavioral Reviews 75:104–113. https://doi.org/10.1016/j.neubiorev.2017.01.039

Louw A, Zimney K, Puentedura EJ, Diener I (2016) The efficacy of pain neuroscience education on musculoskeletal pain: A systematic review of the literature. Physiotherapy Theory and Practice 32:332–355. 10.1080/09593985.2016.1194646

Mack A, Rock I (1998) Inattentional Blindness. MIT Press, Cambridge, MA

Mancini F, Steinitz H, Steckelmacher J, Iannetti GD, Haggard P (2015) Poor judgment of distance between nociceptive stimuli. Cognition 143:41–47. https://doi.org/10.1016/j.cognition.2015.06.004

Martínez M (2011) Imperative Content and the Painfulness of Pain. Phenomenology and the Cognitive Sciences 10:67–90. https://doi.org/10.1007/s11097-010-9172-0

Medoff ZM, Colloca L (2015) Placebo analgesia: understanding the mechanisms. Pain Management 5:89–96. https://doi.org/10.2217/pmt.15.3

Melzack R, Wall PD (1965) Pain mechanisms: A new theory. Science 150:971–979. 10.1126/science.150.3699.971

Melzack R, Wall PD, Ty TC (1982) Acute pain in an emergency clinic: latency of onset and descriptor patterns related to different injuries. Pain 14:33–43. 10.1016/0304-3959(82)90078-1

Merskey H, Albe-Fessard DG, Bonica JJ, Carmon A, Dubner R, Kerr FWL, Lindblom U, Mumford JM, Nathan PW, Noordenbos W, Pagni CA, Renaer MJ, Sternbach RA, Sunderland S

(1979). Pain terms: A list with definitions and notes on usage. Recommended by the IASP Subcommittee on Taxonomy. Pain 6(3):249–252

Moseley GL, Arntz A (2007) The context of a noxious stimulus affects the pain it evokes. Pain 133(1-3):64–71. 10.1016/j.pain.2007.03.002

Paul K, Tik M, Hahn A, Sladky R, Geissberger N, Wirth E-M, Kranz GS, Pfabigan DM, Kraus C, Lanzenberger R, Lamm C, Windischberger C (2021) Give me a pain that I am used to: distinct habituation patterns to painful and non-painful stimulation. Scientific Reports 11:22929. https://doi.org/10.1038/s41598-021-01881-4

Paulson PE, Minoshima S, Morrow TJ, Casey KL (1998) Gender differences in pain perception and patterns of cerebral activation during noxious heat stimulation in humans. Pain 76:223–229. 10.1016/s0304-3959(98)00048-7

Pautz A (2010) Do Theories of Consciousness Rest on a Mistake? Philosophical Issues 20:333–367. https://doi.org/10.1111/j.1533-6077.2010.00189.x

Pautz A (2014) The Real Trouble with Phenomenal Externalism: New Empirical Evidence for a Brain-Based Theory of Consciousness. In: Brown E (ed) Consciousness Inside and Out: Phenomenology, Neuroscience, and the Nature of Experience, Springer, p 237–298

Pereplyotchik D (2017) Pain and consciousness. In: Corns J (ed) The Routledge Handbook of Philosophy of Pain. Routledge, London, p 210–220

Price DD, Barrell JJ (2012) Inner Experience and Neuroscience: Merging Both Perspectives. MIT Press, Cambridge, MA

Purves D, Augustine GJ, Fitzpatrick D, Hall WC, LaMantia A-S, McNamara JO, Williams SM (2004). Neuroscience, 3rd edition. Sinauer Associates, Sunderland, MA

Purves D, Morgenstern Y, Wojtach WT (2015). Perception and Reality: Why a Wholly Empirical Paradigm is Needed to Understand Vision. Frontiers in Systems Neuroscience 9:156. https://doi.org/10.3389/fnsys.2015.00156

Pylyshyn Z (1999) Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. Behavioral and Brain Sciences 22:341–423. 10.1017/s0140525x99002022

Rainville P, Duncan GH, Price DD, Carrier B, Bushnell MC (1997). Pain affect encoded in human anterior cingulate but not somatosensory cortex. Science 277:968–971. 10.1126/science.277.5328.968

Raja SN, Carr DB, Cohen M, Finnerup NB, Flor H, Gibson S, Keefe FJ, Mogil JF, Ringkamp M, Sluka KA, Song X-J, Stevens B, Sullivan MD, Tutelman PR, Ushida T, Vader K (2020). The revised International Association for the Study of Pain definition of pain: concepts, challenges, and compromises. Pain 161:1976–1982. 10.1097/j.pain.0000000000001939

Rosenqvist T (2023a) Color and Competence: A New View of Color Perception. In: Viejo JM, Sanjuán M (eds) Life and Mind. Interdisciplinary Evolution Research, vol 8. Springer, Cham, p 73–103. https://doi.org/10.1007/978-3-031-30304-3_5

Rosenqvist TC (2023b) Seeing *with* color: Psychophysics and the function of color vision. Synthese 202:20. https://doi.org/10.1007/s11229-023-04226-y

Skrzypulec B (2021) Spatial content of painful sensations. Mind & Language 36:554–569. https://doi.org/10.1111/mila.12358

Smith B, Ceusters W, Gildberg LJ, Ohrbach R (2011) Towards an Ontology of Pain. In: Mitsu Okada (ed) Proceedings of the Conference on Ontology and Analytical Metaphysics, February 24-25, 2011. Keio University Press,Tokyo, p 23–36

Thompson E (1995) Colour Vision: A Study in Cognitive Science and the Philosophy of Perception. Routledge, London

Tye M (1995a). Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind. MIT Press, Cambridge, MA

Tye M (1995b) A Representational Theory of Pains and Their Phenomenal Character. Philosophical Perspectives 9:223–239. https://doi.org/10.2307/2214219

Tye M (2002) Representationalism and the Transparency of Experience. NOÛS 36:137–151. https://doi.org/10.1111/1468-0068.00365

Walitt B, Nahin RL, Katz RS, Bergman MJ, Wolfe F (2015) The Prevalence and Characteristics of Fibromyalgia in the 2012 National Health Interview Survey. PLoS One 10:e0138024. https://doi.org/10.1371/journal.pone.0138024

Wall PD (1979) On the relation of injury to pain. Pain 6:253–264. https://doi.org/10.1016/0304-3959(79)90047-2. https://doi.org/10.1016/0304-3959(79)90047-2

Wayand JF, Levin DT, Varakin DA (2005) Inattentional blindness for a noxious multimodal stimulus. American Journal of Psychology 118:339–352. https://doi.org/10.2307/30039070

Wiercioch-Kuzianik K, Bąbel P (2019) Color Hurts. The Effect of Color on Pain Perception. Pain Medicine 20:1955–1962. https://doi.org/10.1093/pm/pny285

Wright L (1973) Functions. The Philosophical Review 82(2):139–168. https://doi.org/10.2307/2183766

Wright L (1976) Teleological Explanations: An Etiological Analysis of Goals and Functions. University of California Press, Berkeley