

PAIN WITHOUT INFERENCE

Laurenz Casser

The University of Sheffield

Forthcoming in *Philosophy and Phenomenological Research*.

Please cite published version.

Abstract: A foundational assumption of contemporary cognitive science is that perceptual processing involves inferential transitions between representational states. However, it remains controversial whether accounts of this kind extend to modalities whose perceptual status is a matter of debate. In particular, it remains controversial whether we should attribute inferential mechanisms to the sensory processing underpinning (human) pain experiences. This paper argues that, contrary to recent proposals in the philosophy and science of pain, pain processing is not mediated by inferential transitions. To this end, I show that standard motivations for inferentialism—including appeal to underdetermination, illusion, cue combination, cognitive penetration, perceptual constancy, and invariance—do not carry over to pain. Instead, I suggest that pain’s sensory processing may be better characterised as an idiosyncratic form of transduction, distinguishing it both from paradigmatic perceptual modalities and canonical transducers.

*Could a machine think?—could it be in pain?—
Well, is the human body to be called such a machine?
It surely comes as close as possible to being such a machine.
(Wittgenstein, *Philosophical Investigations*)*

1. INTRODUCTION

A foundational assumption of contemporary cognitive science is that perception depends upon *unconscious inferences*: that what we see, hear, and feel is—by and large—the result of subpersonal processes which ‘draw conclusions’ about the state of the environment from relevant ‘premises’ (Fodor & Pylyshyn, 1981, Gregory, 1970, Marr, 1982, Palmer, 1999, Ritchie, 2022, Rock, 1983). Some such premises are supplied by sensory organs, and amount to *current* (albeit impoverished and ambiguous) information about the world; other premises are (implicitly) encoded in the architecture of perceptual systems themselves, and amount to *prior* (albeit defeasible) information about what the world is typically like. When we perceive, perceptual systems are seen to literally *infer* the most likely estimate of the perceptual scene on the basis of both current and prior information at their disposal. This theory of perceptual processing is known as *inferentialism* (also referred to as *constructivism* or *cognitivism*).¹

Origins of inferentialism date back at least as far as Ptolemy’s *Optics* (c. 160), which argues that the perception of size requires visual systems to make ‘unnoticed judgements’—a view

¹ As I will go on to clarify shortly, unconscious inference can also be understood in statistical rather than logical terms, replacing talk of ‘premises’ and ‘conclusions’ with that of probability assignments to a hypothesis space (see Section 2).

developed in more detail by the Arab mathematician Alhazen centuries later (c. 1030) and maintained by various authors since (see Hatfield, 2002, for a historical overview). In its modern form, inferentialism is typically attributed to Hermann von Helmholtz (1867), a German physicist and physician, who proposed that perceptual systems ‘construct’ coherent images of the environment in a series of ‘mental adjustments’ or ‘unconscious inferences’. Helmholtz’s approach is now standard within contemporary perceptual psychology and credited with remarkable explanatory power: not only is inferentialism considered to offer the only viable solution to the so-called ‘underdetermination problem’ (of which more later), but it promises to explain a variety of notable perceptual phenomena which would otherwise remain obscure (Rescorla, 2015a, b, 2020, Ritchie, 2022).

While inferentialism is frequently discussed in the context of *visual* perception, it is typically understood as a theory about perception quite generally. Consequently, the processes which underlie hearing, touch, and proprioception are similarly seen to implicate inferential mechanisms (Burge, 2010, Palmer, 1999). However, whether or not we should attribute mechanisms of this kind to modalities whose perceptual status is more contentious (such as olfaction and gustation) or, indeed, to sensory systems quite generally, remains a controversial (albeit neglected) area of inquiry.

According to an increasingly influential view, there is compelling evidence in support of inferential processing in (human) pain perception (Seymour & Mancini, 2020, Tabor et al., 2017, Casser & Clarke, 2023). It is argued that pain systems, much like visual and auditory systems, *inferentially* interpret their inputs and thereby determine the state (and causes) of bodily damage incurred by the organism. This view is not unmotivated since, according to its proponents, pain processing faces (and solves) some of the same ‘problems of inference’ as afflict paradigmatic perceptual systems. In fact, this much is taken for granted in recent discussions of pain’s cognitive architecture (Gligorov, 2017, Jacobson, 2017, Shevlin & Friesen, 2020, Skrzypulec, *forthcoming*), as well as increasingly popular Bayesian and predictive coding models of pain (Wager, 2004, Watson et al., 2006, Brown et al., 2008a, 2008b, Seymour and Dolan, 2013, Yoshida et al., 2013, Moutoussis et al., 2014, Anchisi & Zanon, 2015, Wiech, 2014, 2016, Ongaro et al., 2019, Hoskin et al., 2019).

These developments are not inconsequential: to this day, traditional pain science seeks to explain how pain systems operate at a primarily neurophysiological and functional level of analysis. However, if an adequate explanation of pain processing requires the ascription of *inferences* (and, by extension, *mental representations*), then this suggests that there are generalisations which can only be captured at a more abstract, ‘symbolic’ or ‘computational’ level of description (see Pylyshyn, 1984). Consequently, if pain *is* inferential, then a complete science of pain must extend its current vocabulary to capture such generalisations—a trajectory reminiscent of that of vision science in the 1970s (see Marr, 1982).

However, contrary to recent proposals in the philosophy and science of pain, this paper argues that pain processing is *not* inferential. To this end, I show that standard motivations for inferentialism—including appeal to underdetermination, illusion, cue combination, cognitive penetration, perceptual constancy, and invariance—do not successfully carry over to pain. More specifically, I argue that (i) pain processing does not solve an underdetermination problem (rendering the introduction of inferential mechanisms unmotivated), that (ii) candidate cases of pain illusions, cue combination, and cognitive penetration fail to successfully motivate a problem of inference for pain (though they may at times motivate a problem of inference for other sensory

systems), and that (iii) potential further indicators of inferential underpinnings, such as constancy mechanisms and invariance, are absent from pain processing (and hence needn't be accounted for inferentially or otherwise). As a consequence, pain does not manifest any of the prominent explananda which inferentialism is intended to explain, rendering the introduction of inferential mechanisms gratuitous and, hence, unjustified. I conclude by sketching an alternative model of pain's sensory processing as an idiosyncratic form of transduction, which differentiates pain both from paradigmatic perceptual modalities and canonical sensory transducers.

This paper is divided into four sections. Section 2 provides an overview of inferentialism in perceptual psychology and outlines why perceptual systems are seen to make inferences. Section 3 traces recent commitments to inferentialism in the philosophy and science of pain, and argues that despite apparent similarities between pain and paradigmatic forms of perception, none of the standard motivations for inferentialism carry over to pain. In section 4, I gesture at an alternative conception of pain processing as a form of sensory transduction.

2. MAKING AN INFERENCE

It is said that perception is all about the news. Its purpose is to gather information about the current state of the organism's environment—the 'here and now', as it were—which will assist the organism in the conduct of its life (Armstrong, 1968: 209; Block, 2014: 567, 2022: 99, cf. Phillips & Firestone, 2023). To do so, perceptual systems receive inputs via specialised sensory receptors and form (more or less) accurate representations of distal objects, scenes, and events on their basis. This process is not straightforward. Specifically so, since said inputs typically *underdetermine* the state of the world: one and the same pattern of proximal sensory stimulation is compatible with multiple (sometimes infinitely many) possible distal causes.

To get an intuitive grasp on this issue, consider touch: When we explore objects with our hands, we typically have a 'feel' for the unity of objects we are touching (see Lederman & Jones, 2011). When we grasp an orange, for example, we *feel* that we are in contact with a single, contiguous object. When we are clutching a bunch of blueberries, on the other hand, we *feel* that we are in contact with several, disconnected items. Notably, however, in the absence of visual aid, the only proximal stimulation our tactile system has to go on when estimating unity or disunity is information from separate contact points of our skin: information to the effect that there is something of a certain texture, temperature, and force applied 'here', 'here', and 'here'.² That information, however, is compatible with multiple interpretations: it may very well be the case that there is only one contiguous object that is responsible for all points of contact, but it might also be the case that there are multiple objects that are. Some such interpretations may, of course, be more plausible than others. The point, however, is simply that the sensory inputs which the system is receiving are themselves ambiguous between these distinct possibilities. One and the same pattern of stimulation is compatible with multiple ways the world might be. This is the *underdetermination problem*.

²This is an oversimplification. Even in the absence of visual information, the tactile system receives additional inputs from proprioception, which help it try and determine relevant properties of objects, such as shape (Hsiao, 2008).

Underdetermination is a pervasive phenomenon, affecting all perceptual modalities in multiple ways.³ Remarkably, however, perceivers tend to have a pretty good idea of the world around them—underdetermination notwithstanding. Indeed, our senses typically provide us with stable and fairly accurate information about the state of our environment, including information about the number of objects we are grasping. How is this possible? The classical solution proposed by Helmholtz has been that perceptual systems interpret their inputs on the basis of ‘implicit assumptions’ about environmental regularities and *infer* the most likely distal conditions. For instance, when the tactile system is trying to determine the number of objects in contact with the perceiver’s hand, it may rely on information about the kinds of stimulation patterns that objects in the perceiver’s environment are likely to cause. And so it may assume, for example, that ‘contiguous objects are fairly uniform in temperature’, that ‘the likelihood of contact with multiple objects increases with the spatial separation between contact points’, or that ‘single, contiguous objects are unlikely to touch the “outsides” of two adjacent fingers simultaneously’. Assumptions of this kind allow the system to rule out a number of possible interpretations of its inputs and to ultimately make an informed estimate of the quantity of objects in contact with the perceiver.

To give a concrete (albeit idealised) example of such a case, imagine that the tactile system is receiving inputs from mechanoreceptors located on the left of the subject’s right index finger *and* inputs from mechanoreceptors located on the right of the subject’s right middle finger (the ‘outsides’ of those fingers, if you will). Here, the tactile system is faced with inputs that are ambiguous between at least two interpretations: Either one contiguous object is in contact with the outsides of middle and index finger, or two objects are in contact with the outsides of middle and index finger respectively. Since contiguous objects in the perceiver’s environment are typically shaped in such a way that they do not simultaneously come into contact with the ‘outsides’ of two adjacent fingers, the tactile system may assume, quite generally, that ‘single, contiguous objects are unlikely to touch the outsides of two adjacent fingers simultaneously’. On the basis of this assumption, the system is then able to infer that the former interpretation is far less likely to be correct than the latter, leading it to conclude that the perceiver is in contact with multiple objects. In this way, the system can be seen to have determined the state of the environment by way of inference.

Perceivers may, of course, encounter the world under *atypical* conditions, in which case the system’s implicit assumptions are likely to be violated. A perceiver may, for example, cross their middle and index finger and hence touch a single object with the ‘outsides’ of two adjacent fingers after all. In this case, the system’s implicit assumptions will lead it to *erroneously* conclude that the perceiver is in contact with multiple objects—an effect known as *Aristotle’s Illusion* or *tactile diplopia* (Baysan & McPherson, 2017). However, given the assumptions that the tactile system is hypothesised to be making, this is just as a proponent of inferentialism would predict. In fact, the inferentialist can quite naturally account for perceptual illusions more broadly as instances of inferential failure (Fodor & Pylyshyn, 1981).

As these examples illustrate, the notion of inference not only helps explain how perceptual systems get things *right*, even in the face of underdetermination, but also how perceptual systems get things *wrong*, as in the case of perceptual illusions. In addition, an inferential

³ Examples of underdetermination abound. For visual cases see Burge (2010), Palmer (1999), Ramachandran (1988). For auditory cases see Nudds (2015).

understanding of perceptual processing is seen to explain a variety of other notable perceptual mechanisms and phenomena (see Section 3), including cue combination (Rescorla, 2020), (alleged) cognitive penetration (Hohwy, 2013), perceptual constancy (Rescorla, 2015a), and invariance (Ritchie, 2022).

Since its introduction by Helmholtz, inferentialism has been developed in roughly two different ways. On one version of the view, the notion of inference is construed in something like a logical sense: here, the perceptual system's 'implicit assumptions' are regarded as stored premises which feature in deductive, inductive, or abductive lines of reasoning. Inferentialism, in this sense, posits that perceptual systems 'draw conclusions' from relevant 'premises' in a rather literal sense (see e.g. Rock, 1983: 272-282). On another version of the view, the notion of inference is construed in a statistical, Bayesian sense: here, talk of 'implicit assumptions' is replaced with that of a hypothesis space reflecting possible environmental conditions to which probabilities are assigned. Upon receiving incoming sensory evidence, these probabilities are reallocated in (rough) accordance with Bayes' Law, and a favoured hypothesis is selected. Inferentialism, in this sense, posits that perceptual systems approximately conform to norms of Bayesian inference (see e.g. Rescorla, 2015a). While these two versions of inferentialism differ from one another in significant respects, which version one endorses does not bear on the arguments presented in this paper.

It is important to note that inferentialism of either strand tends to take the role of inferences in perceptual processing quite literally: perceptual systems don't operate *as if* they performed inferences, but *in fact* do so. However, since said inferences are unconscious, and hence inaccessible to the perceiver, their presence cannot be verified introspectively. Rather, their postulation is justified, on the inferentialist's view, *only* by the fact that a description of perceptual processing without appeal to inferential mechanisms falls short of adequately explaining the perceptual phenomena in question. This approach is seen to be vindicated by the fact that inferentialism has led to major advances in the recent history of perceptual psychology.

What exactly (unconscious) inferences are supposed to be is not always so clear, and depends on the species of inferentialism one endorses. However, proponents of either version of inferentialism typically understand inferences to be paradigmatically *non-associative* and *reason-responsive* transitions between representational states (see Casser & Clarke, 2023). Inferences are 'non-associative' insofar as relevant transitions are not (merely) linked by an associative relation, such as 'closeness' or 'similarity', which can be established through conditioning. Rather, they are guided by a set of rules or principles, such as, e.g., disjunctive syllogism or Bayes' Law, which convey an internal logic to perceptual processing. Relatedly, inferential transitions are 'reason-responsive' insofar as they can be modified by evidence, such as, for example, content-relevant information from another sensory system (see Jenkin, 2022).

Since inferences in perceptual processing involve transitions between contentful states, such as premises or estimates *about* possible environmental conditions, unconscious inferences are seen to implicate *mental representations* (see Fodor & Pylyshyn, 1981, Rescorla, 2015a). Mental representations are 'mental' insofar as they are proprietary to psychology, and 'representations' insofar as they are 'causally active semantic mediators' standing in for something other than themselves, such as the state of the environment or an estimate thereof.

Inferentialism is, of course, not the only available theoretical framework by which perceptual processing may be interpreted. So-called 'ecological' or 'embedded' accounts of perception, for example, remain competitive, albeit marginal, alternatives (see Gibson, 1979,

Orlandi, 2014). In these accounts, perception is not mediated by inferences or mental representations, but capable of ‘picking up’ environmental properties *directly*. Accordingly, what is seen to explain the perceptual solution to the underdetermination problem, for example, is not that perceptual systems somehow *represent* certain environmental regularities and make *inferences* on their basis, but that the environmental regularities themselves have shaped perceptual systems in the course of evolution to process ambiguous sensory stimulation one way rather than another. As such, perceptual processing is seen to be explained by ‘the facts’ as opposed to perceptual systems’ representation of the facts (Orlandi, 2014). However, ecological and embedded accounts remain marginal as they are widely seen to have (insurmountable) difficulties accounting for perceptual illusions, constancies, and adaptation effects (Rescorla, 2015a, b).

For the purposes of this paper, I will assume (primarily for the sake of argument) that such criticisms of ecological and embedded accounts are decisive, and that inferentialism is the most promising account of perceptual processing. In fact, I am going to assume that the phenomena of underdetermination, illusion, cue combination, cognitive penetration, perceptual constancy, and invariance are directly indicative of inferential underpinnings to sensory processing. I will argue, however, that even if this much is granted, inferentialism does not carry over to pain.

3. INFERRING PAIN

The science of perception and the science of pain have evolved fairly independently from one another (see Finlay, 2019). And whereas inferentialism has been prevalent in the science of perception for some time now, it remains a rather recent proposal in the philosophical and scientific study of pain. Putting these differences aside, however, a central motivation for inferentialism about pain is that pain systems are seen to face similar processing tasks as do paradigmatic perceptual systems—tasks which require the system to make inferences about the state of its (bodily) environment in order to succeed. For example, Casser and Clarke (2023) suggest that:

[P]ain may be regarded as inferential *in much the way vision is*. For just as a retinal image underdetermines the size and shape of the objects one perceives, leaving these to be inferred by visual mechanisms, somatosensory inputs to polymodal nociceptors may underdetermine the nature of bodily damage (see Ringkamp et al., 2013). By parity of reasoning, this suggests that for pain processes (of the sort posited by perceptualists) to reliably identify objective levels of bodily damage, they, too, perform inferences akin to those performed by other perceptual mechanisms (p. 3, emphasis added).⁴

Similarly, Tabor and colleagues (2017) write that

There is now a compelling body of literature to support an inferential model of pain. During the experience of pain, *just as during other perceptual experience*, the brain makes

⁴ As I will elaborate on in Section 3.1.1., Casser and Clarke ultimately propose that pain’s computational function diverges from that of vision insofar as it makes further ‘practical inferences’ based on the state of bodily damage already inferred.

inferences based on incomplete information. Specifically, the most common trigger of pain is a somatosensory barrage that includes, but is not limited to, activity in high threshold primary receptors (nociceptors) and their projections. Physiologically, nociceptive input is always accompanied by—indeed, preceded by—a wide array of non-nociceptive input triggered by other somatosensory receptors and a multisensory suite of event-related information. This suite of information needs to be integrated with prior knowledge and over time, in order to calculate the experience that would most favourably serve the immediate objectives of the organism (p.3, emphasis added).

As these authors suggest, pain faces similar processing tasks as do paradigmatic perceptual systems, and hence inherits similar (computational) problems. Since plausible solutions to these problems are widely seen to implicate inferential operations, pain processing is supposed to be inferential in much the same way perceptual processing is standardly taken to be—whether pain is itself a perceptual phenomenon (strictly speaking) or not.

3.1. Underdetermination

In typical cases of underdetermination, there is a one-to-many relation between sensory input and possible distal cause: one and the same pattern of stimulation is compatible with multiple ways the world might be. In order to determine which way the world actually is, perceptual systems are seen to perform genuine inferences. Casser and Clarke (2023) suggest that there are at least two ways of motivating an underdetermination problem for pain, each of which is broadly aligned with one of the two most prominent views on pain in contemporary philosophy: *perceptualism* and *imperativism*.

3.1.1 Two Kinds of Underdetermination

According to perceptualists, pain is a perceptual phenomenon akin to paradigmatic forms of perception, such as vision and audition. On this view, pain resembles other perceptual modalities insofar as it functions to inform organisms about objective states of their physical environment—in this case: damage or threat of damage to the organism's body (Armstrong, 1968, Pitcher, 1970). In doing so, pain plausibly informs not only about the existence of such damage, but also about its location, its extent, and (perhaps) even its immediate causal origins (as represented by different pain qualities; see Tye, 1995a, b).

If perceptualism is broadly correct, Casser and Clarke suggest, then it is natural to expect that pain faces an underdetermination problem akin to that of other perceptual modalities (p. 3f.). After all, if we believe perceptualists, pain *is* a perceptual modality. And if other forms of perception cannot uniquely determine the state of the organism's environment on the basis of proximal sensory stimulation, then we might expect the same to hold for pain. And indeed, there is some plausibility to this. For many of the peripheral receptors involved in pain processing (so-called *nociceptors*) are polymodal: they respond to multiple stimulus modalities, including mechanical, thermal, and chemical stimuli. On the assumption that pain systems are interested in

telling the difference between these, additional information will be required in order to determine the stimulus-modality and, with that, the nature of bodily damage incurred.

Of course, perceptualism is controversial, and not everyone agrees that pain shares the functional role of perception (Coninx, 2019, Casser, 2021, Rosenqvist, *forthcoming*). Imperativists, for example, reject perceptualism and argue that pain is better characterised as a form of behavioural homeostasis, akin to hunger or thirst (Martinez, 2011, Klein 2007, 2015, Barlassina & Hayward, 2019). In their view, pain is not primarily concerned with informing the organism about the state of its environment, but with promoting suitable protective behaviours (such as retraction of a limb or nursing of a body part) in order to avoid serious injury and subsequent death. Critically, if imperativism is to be preferred over perceptualism, then we cannot straightforwardly infer an underdetermination problem from pain's alleged kinship to perception, as imperativism denies that such kinship exists.

However, Casser and Clarke suggest that a broadly imperativist view of pain may motivate an underdetermination problem of its own. One consideration is that a system concerned with promoting adequate protective behaviours may still need to determine the organism's state of bodily damage, since the effectiveness of a protective response will depend, at least in part, on facts about what the organism needs protection *from*. And since the determination of such facts is likely to involve underdetermination, as we said before, an imperativist conception of pain will similarly suggest an underdetermination problem. A further consideration is that there might be *additional* underdetermination at the level of the behavioural response being promoted, since one and the same bodily state might warrant distinct protective responses, depending on the behavioural context. For example, incurring a serious bone fracture in a safe and familiar environment might best be addressed by immediate pain-onset to facilitate nursing behaviours which keep the affected area from further strain; however, the very same injury incurred in a fight or flight situation, such as on the battlefield, might best be addressed by sufficiently *delayed* pain-onset, which will allow the organism to deal more effectively with the dangerous situation at hand. In this way, pain processing might be facing underdetermination twice over: first at the level of determining the state of bodily damage on the basis of proximal sensory stimulation, and second at the level of determining a suitable protective response on the basis of the state of bodily damage inferred.

If Casser and Clarke are correct, then there exists plausible reason to expect pain processing to be faced with an underdetermination problem irrespective of whether one ultimately prefers a perceptualist or an imperativist perspective on pain. And since underdetermination is generally regarded as providing some of the most compelling evidence for inferentialism in the perceptual domain, parity of reasoning suggests that we should regard underdetermination as providing similarly compelling evidence in the case of pain.

3.1.2. From Underdetermination to Inference

Pace Casser and Clarke, I want to suggest that the argument from underdetermination outlined above is not as plausible as these authors take it to be. To illustrate why, it is helpful to consider some of the individual claims that their argument relies upon in order to yield the desired conclusion that pain processing involves inferences:

1. First and foremost, it relies on the claim that pain processing is indeed facing underdetermination to begin with. If pain processing did not face underdetermination, there would obviously be no ‘argument from underdetermination’.
2. Second, supposing that pain processing *is* facing underdetermination, it relies on the further claim that this poses a *problem* for the system. If the presence of underdetermination did not pose a problem, there would be no reason to expect an inferential solution, since there wouldn’t be anything that needs to be solved.
3. Third, assuming that pain processing is facing underdetermination *and* that this poses a problem for the system, it relies on a yet further claim to the effect that the system *overcomes* or *solves* this problem. If it did not solve the problem, there would be no reason to appeal to inferences in order to explain its solution.

In order for there to be an argument from underdetermination and for that argument to yield the conclusion that pain processing involves inferential mechanisms, each of these claims has to be accepted. In the case of paradigmatic perceptual systems, they typically are (*mutatis mutandis*): virtually no one denies that the visual system, for example, faces underdetermination, that this is a problem, and that the system solves it. In fact, even opponents of inferentialism typically admit this much. In the case of pain, however, I suggest that there are compelling reasons to be sceptical of each of these claims.

1. *Is pain processing really facing underdetermination?* There are at least two reasons to be doubtful of an affirmative response. The first is that the causal determinants of pain experiences (unlike the causal determinants of e.g. visual or auditory experiences), are not very well understood. In fact, unlike vision science, for example, which has a very detailed understanding of the precise causes of visual states, and which can boast rigorous mathematical models of underdetermination as a result, pain science remains in *desperate* search for the causes of pain. In fact, many kinds of pain, including migraines, headaches, lower back pain, and so forth, famously occur without any accurately discernable causes (Melzack & Wall, 2008). Moreover, it has become increasingly clear that pain experiences are influenced by a wide variety of different psychological states and events which go far beyond mere nociceptive input. I take it that this is (at least in part) why pain medicine and pain management remain as imperfect as they are: not enough is known about the causes and determinants of pain to reliably prevent it. However, if we know little about how pain states are determined, then it is unclear how we could know that such states are *under-*determined.

The second reason to be sceptical is that the percepts of pain experiences, ‘pains’, appear far less determinate and complex than the percepts of other sensory modalities. Pains differ in terms of their sensory qualities, bodily location, and intensity, but they plausibly lack the complexity of visual, auditory, and tactile objects, which appear to us in considerable detail. This observation is not inconsequential, since it illustrates a difference in how demanding we may expect the processing tasks of these systems to be. For example, part of what makes the work of the visual system so difficult is that it needs to create three-dimensional percepts from two-dimensional retinal stimulation (Burge, 2010). However, pains are typically not experienced as three-dimensionally extended, nor do they have distinct shapes, sizes, textures, or any other comparably complex properties which would indicate that the system needs to contribute

additional information to interpret its inputs. The same type of worry applies to an imperativist view on which what is underdetermined are not the proximal causes of pain, but the protective behaviours pain is seen to promote. Considering that the number of candidate behaviours (retraction, nursing, rest) is fairly limited, it is unclear why we should think that the system couldn't determine the appropriate output simply on the basis of information it receives from its inputs. To return to the case of bone fracture, for example, it is unclear why the system's decision between immediate and delayed pain onset couldn't be determined solely by, e.g., nociceptive input plus the organism's adrenaline levels or heart rate (standing in as proxies for the behavioural context), hence leaving no need for inferences. The critical point, then, is that whereas underdetermination in the perceptual realm has long been known as scientific fact, ignorance of which may have you fail science class, underdetermination in the case of pain remains a mere desideratum for the inferentialist.

2. *Is underdetermination really a problem for pain systems?* But let us assume for the sake of argument that there *is* underdetermination: that the proximal sensory stimulation pain systems receive as inputs underdetermines that state of bodily damage, or that the state of bodily damage underdetermines candidate protective behaviours, or both. Even if this much is granted, neither kind of underdetermination will automatically be a *problem* for the system unless we make fairly specific assumptions about the system's tasks and goals. For example, if we suppose that nociceptive stimulation plausibly underdetermines whether an injury to the body is thermally or mechanically induced, then this will be a problem for the system only on the further assumption that it is among the system's goals to determine which of these alternatives it is. On a perceptualist view of pain, this assumption may seem quite natural. However, it is worth noting that if pain systems do aim at distinguishing between thermally and mechanically induced injuries in the formation of their percepts, they do not seem to do so very successfully. In fact, subjects generally fail to tell apart thermally from mechanically induced pains, even if explicitly instructed to do so in a laboratory setting (see Wall & McMahon, 1986). Nor is it clear that such differentiation would serve any meaningful purpose on an imperativist view. Whether an acute injury is thermally or mechanically induced hardly matters for imminent behavioural purposes: retraction of the limb and nursing of the affected body part would seem to be an appropriate response in either case. But if that is true, it remains unclear whether underdetermination, even if it does affect pain systems, constitutes a problem for such systems (irrespective of whether we are leaning towards a perceptualist or imperativist conception of pain).

3. *Do pain systems really overcome an underdetermination problem?* But now, let us assume further (and again, for the sake of argument), that pain systems not only face underdetermination, but that underdetermination also poses a problem for them. I'd like to suggest that even if this much is granted (and I think it's granting a lot), there is little reason to suppose that pain systems actually *solve* this problem. After all, the correlation between pain and bodily damage is notoriously tenuous, as is the correlation between pain intensity and physical trauma, as well as between pain quality and stimulus modality. Unlike paradigmatic perceptual modalities, which keep us remarkably well informed despite underdetermination, pain frequently leaves us in the dark (Wall, 1979). And so, even if pain faces an underdetermination problem akin to that of paradigmatic perceptual systems, it is a stretch to say that it is 'solving' it. Nor is a solution to underdetermination more plausible on an imperativist view: bodily trauma frequently occurs without any immediate onset of pain, even in behavioural contexts in which pain onset would be helpful (since it would encourage timely treatment). As Melzack and colleagues (1982) found in

a large-scale study, 37% of emergency room patients did not feel pain at the time of injury despite the fact that many of them incurred their injuries under non-threatening conditions at home or at the workplace. If pain systems face an underdetermination problem at the level of behavioural response, aiming to induce pain in situations in which pain behaviours are helpful and to prevent pain in situations in which pain behaviours are distracting, they hardly provide an effective solution to it.

As these considerations illustrate, an argument for inferentialism about pain on the basis of underdetermination faces serious difficulties at each individual stage. Not only is it yet to be shown that pain systems actually face underdetermination to begin with, but it remains unclear whether underdetermination would even constitute a problem for pain systems, much less a problem they are equipped to solve. In the absence of concrete evidence and further argumentation, underdetermination in pain processing remains no more than a desideratum for inferentialists.

3.2. Illusion

Illusions are a familiar type of perceptual experience in which subjects perceive an object but misperceive one or more of its properties. In cases of this sort, perceptual objects *appear* to be some way other than they *really* are. Inferential accounts of perceptual processing are regarded as particularly well equipped to explain such cases insofar as they can identify their ‘sources’: namely, the relevant perceptual system’s prior assumptions that are being violated (Rescorla, 2015a).

However, in the case of pain, a distinction between appearance and reality is controversial. Indeed, a popular view holds that pain simply *is* the way it is *felt* to be, and hence remains immune to the possibility of illusion and hallucination. On this view, inferentialism about pain couldn’t be motivated by appeal to its explanatory power of illusions, simply because there are no illusions to be explained. Having said that, everyone must contend with the fact that there are a number of phenomena which are (rightly or wrongly) recognised as candidate cases of ‘pain illusions’—including phantom limb pain, referred pain, and the so-called ‘thermal grill illusion’—and which pose an explanatory challenge. As such, it is worth considering whether the introduction of inferential mechanisms is required for this challenge to be met.

Perhaps the most promising candidate of a putative pain illusion indicating inferential underpinnings is the thermal grill illusion. In the TGI, thermal stimuli alternating in temperature between innocuously warm and innocuously cold are applied to the subject’s skin. Even though subjects do not experience these thermal stimuli as painful when applied individually, they experience a burning pain sensation when applied collectively—provided the stimuli are suitably arranged and their temperature set to differ by at least 20°C (Thunberg, 1896). Subjects’ pain is said to be ‘illusory’ in the (somewhat idiosyncratic) sense that it is generated by an entirely innocuous physiological event. Touching the thermal grill ‘shouldn’t’ be painful, but it is.⁵

The question of why the TGI occurs remains unresolved (see Bouhassira et al., 2005, Green, 2002, Craig and Bushnell, 1994, Fardo et al., 2019). However, the phenomenon might be

⁵ Interestingly, the painfulness of the TGI is not entirely uncontroversial (see e.g. Bach et al., 2005). Historically, the TGI was regarded as a phenomenon of synthetic heat from its discovery in 1896 until the 1990s, and only then became known as a pain phenomenon. However, for the purposes of this discussion, I will put these difficulties aside.

described as involving a kind of ‘mistake’ in the pain system’s identification of the stimulus. One way of understanding this process is as a matter of inferential failure: the relevant pain systems interpret incoming sensory information by drawing on certain inbuilt assumptions about environmental regularities. In the case of the TGI, these inbuilt assumptions are (somehow) violated, leading the system to perform a bad inference and conclude that the stimulus is to be deemed painful for the organism. What exactly the implicit assumptions relevant to the TGI are meant to be (provided there are such assumptions) remains speculative, but one might think that the system expects innocuous objects in contact with the body to be fairly uniform in temperature. If, however, as in the case of the TGI, juxtaposed contact points differ in temperature by a sufficiently large margin, the system (erroneously) judges that the stimulus threatens tissue damage and should be avoided.

Unfortunately, however, this suggestion faces serious difficulties. We said that the TGI intuitively involved a type of ‘error’ or ‘mistake’ in the sense that the thermal stimuli *shouldn’t* be giving rise to a pain experience (as they are evidently innocuous). This error, we said, might helpfully be cashed out in terms of inferential failure. However, it is important to note that whatever error or inferential failure the TGI may or may not suggest, it is at best unclear why we should think this error to lie at the level of *pain* processing, as opposed to processing that may precede it. That is because the ‘mistake’ involved in the TGI is not that the experience presents the pain in some way it is not, or fails to present the pain in some way it is, but rather that there is a pain experience *at all*. And so, unless we thought that it is somehow among the processing tasks of a dedicated pain system to work out itself whether or not a given stimulus should lead to a pain experience—which is not among the tasks we ordinarily attribute to sensory systems—the TGI doesn’t actually suggest any error at the level of pain processing itself. Rather, the TGI might be interpreted as involving, say, an (inferential) failure at the level of thermal or tactile processing, leading the relevant systems to erroneously conclude that the stimulus is considerably hotter than it actually is—a conclusion which triggers pain, but which isn’t the result of a pain system’s erroneous processing.

This alternative possibility is not unmotivated, as can be illustrated by consideration of the phenomenon of *thermal referral* (see Lederman & Jones, 2011). The setup of such cases is fairly similar to that of the TGI: the ring and index finger of one hand are both placed on a warm or cold thermal stimulator, while the middle finger is situated on a neutral surface. Upon contact, subjects have been found to perceive the temperature of the neutral surface as having the same temperature as the thermal stimulators under their adjacent fingers. Moreover, it has been found that such thermal stimulation of adjacent fingers can in fact *enhance* the perceived warmth or cold of a surface touched by the middle finger (Green, 1977). As these effects suggest, the thermal inputs to adjacent sites of the body are not processed independently. Instead, thermal processing appears to take into account the temperatures of neighbouring stimuli, which influences the perceived temperature of the surface.

The relevance of this finding is that it provides evidence of the fact that tactile processing involves the estimation of stimulus temperature, and that it may mistakenly conclude that said stimulus temperature is (significantly) higher than it actually is. This result opens the door for the following hypothesis: namely, that the TGI involves a mistaken temperature estimate at the level of tactile processing, the output of which (i.e. a signal of dangerously high stimulus temperature) then causes a pain experience. Critically, on this picture, the only stage of (inferential) stimulus interpretation is that of tactile processing, whereas the pain system simply takes as its input the

very temperature estimate already worked out beforehand. If we accept this interpretation of events, the TGI fails to motivate a problem of inference for pain, though it may be seen as motivating a problem of inference for touch.

Unless proponents of inferentialism can show that this proposed interpretation of the TGI is untenable, or, alternatively, identify other illusions which demonstrate the need for inferential mechanisms in pain processing more clearly, candidate cases of pain illusions like the TGI do not provide evidence for inferentialism about pain.

3.3. Cue Combination

Organisms frequently receive information (or “cues”) about a single environmental variable from multiple sources. The perceived *size* of an object, for example, may derive from both visual and haptic information; the perceived *location* of an object may derive from both visual and auditory cues, and so forth. In cases of this kind, perceptual systems integrate distinct sensory cues to aid their analyses and frequently produce more accurate estimates of environmental properties as a result. This phenomenon is known as *cue combination*.

As is frequently stressed by proponents of inferentialism, cases of cue combination are naturally explained by appeal to unconscious inferences (Rescorla, 2020). On this view, perceptual systems form hypotheses about environmental properties under uncertainty, and take on new representational content derived from other sources in non-associative, reason-respecting ways; a process which leads the relevant systems to draw a (more) accurate conclusion. Bayesian perceptual frameworks, which lean heavily on the notion of inferential processing, are seen to model sensory integration of this kind with particularly great success (see Trommershäuser et al., 2011).

Proponents of Bayesianism have identified a variety of potential cue combination effects involving pain. However, it is important to note that many such cases do not meet the criteria for *inferential* processing by the standards of mainstream inferentialists. For example, Tabor and colleagues (2017) discuss an experiment by Moseley and Arntz (2007) in which the experimenters paired nociceptive with visual cues. Subjects would receive a noxious stimulus to their hand and be shown a red or blue visual cue, after which they would rate the stimulus in terms of temperature, pain unpleasantness, pain intensity, and pressure, on a visual analogue scale. As the experimenters found, noxious stimuli paired with a red cue are rated hotter and more painful than noxious stimuli paired with a blue cue. On the assumption that subjects associate red light with heat and danger, and blue light with cool and safety, Tabor and colleagues propose that nociceptive and visual cues are combined to estimate the level of threat to the subject’s body.

However, pace Tabor and colleagues, it is difficult to accept Moseley and Arntz’s finding as evidence for inferential processing. One reason for this is that it is unclear from the experimental results if the detected effect pertains to pain processing itself or to the subjects’ judgement of pain. In other words, it is unclear whether subjects rated the noxious stimuli paired with red light as hotter and more painful than those paired with blue light because these stimuli really did *feel* hotter and more painful, or because subjects merely judged them to be so. The former possibility would indicate an effect on pain processing (as Tabor and colleagues are assuming it is), the latter merely a kind of response bias. However, subject ratings alone cannot

tell these alternatives apart. Hence, it is unclear whether the phenomenon in question even constitutes a genuine case of cue combination or not.

Another reason to be sceptical of Tabor and colleagues' interpretation of this study is that the measured effect (even if it is an effect on pain processing) does not appear non-associative and reason-respecting as is ordinarily required for genuine inferences (see Section 2). After all, the relation between 'the colour red' and 'more pain' or 'the colour blue' and 'less pain' is hardly one of content. If pain processing did indeed take on a visual cue to the effect that there is red or blue light in the vicinity, that cue would itself provide no *reason* for the system to revise its analysis and increase or decrease the intensity of pain. And so, if there is a cue combination effect that affects pain processing here, that effect appears associative and not inferential.

A perhaps more compelling case of cue combination has been identified by Casser and Clarke (2023). These authors draw attention to cases in which subjects report an analgesic effect as a consequence of visual feedback, especially mirror therapy treatment of phantom limb pain. As part of this treatment, patients with phantom pain view a mirror reflection of their intact limb at the location at which their phantom limb is felt to be, thereby creating an illusory percept *as of* an undamaged body part. This experience is reported to have an analgesic effect (see Ramachandran & Altschuler, 2009). On one possible interpretation, the effect is due to an integration of nociceptive information (signalling bodily damage at the location of the phantom limb) and visual information (signalling no visible damage at the location of the phantom limb). As the visual information contradicts the nociceptive information, the relevant pain systems lower their confidence in bodily damage at the location of the phantom limb upon integration, leading to an analgesic effect. This, Casser and Clarke suggest, has the appearance of an inferential transition.

However, there are at least two problems with this particular suggestion. The first problem is that there is a question of how robust the data for mirror therapy-induced analgesia really is. For despite the fact that the relevant studies cited above record a change in subjects' estimation of pain intensity, and despite the fact that mirror therapy has become a recognised treatment option for phantom pain, a number of larger clinical trials and textbook entries on mirror therapy find no significant analgesic effect during or after such treatment (Brodie et al., 2007, Nikolajson, 2013); a finding that is possibly reflected in the observation that mirror therapy is used significantly less frequently than are its alternatives—such as pharmacological treatment and physical therapy—even though it is inexpensive and free of known side effects (Nikolajson, 2013). There is, as such, a question as to whether or not the phenomenon which inferentialism is seen apt to explain reliably occurs.

The second problem resembles that encountered in the context of the TGI (see Section 3.2): namely, that even if said effects did reliably occur, and even if they were indicative of inferential underpinnings, it is unclear why we should attribute them to the processing of a pain system. To see why this is not a given, consider that amputees frequently report that their phantom limb contorts into an extremely painful spasm upon amputation, which feels as if the patient's nails are 'digging into their palm' (Ramachandran et al., 1995: 489, Ramachandran & Altschuler, 2009: 1697). In the mirror condition, however, amputees can typically (and often for the very first time) feel their phantom limb obey their motor commands and move into a relaxed position, reducing their pain (ibid.). This suggests that a patient's phantom pain is (at least in part) a consequence of illusory limb position, and that the mirror therapy condition involves an effect on proprioceptive processing. If so, we can formulate an alternative to Casser and Clarke's

interpretation of the situation: It's not that a pain system (inferentially) integrates visual information which leads it to a new conclusion about the state of bodily damage, but rather that a proprioceptive system (inferentially) integrates visual information which leads it to revise its estimation of limb position. And whereas its former estimate as of a distorted body part was causing a pain experience, its new and revised estimate does not. On this picture, the only stage of inferential processing is that of proprioceptive (and visual) processing, whereas the pain system simply takes as its input the very estimate already worked out beforehand. If we accept this interpretation, then mirror therapy fails to motivate a problem of inference for pain (similarly to the TGI), though it may be seen as motivating a problem of inference for proprioception.

This is not to say that all phantom pain is guaranteed to be explained in terms of proprioceptive analysis, nor, of course, that there couldn't be other cases of cue combination which would provide more compelling evidence for inferential underpinnings in pain processing. However, if there are, then it is on the proponent of inferentialism to identify and motivate such cases. For, as we have seen, the candidate cases of cue combination under discussion do not provide evidence for inferentialism about pain.

3.4. Cognitive Penetration

Similarly to cases of cue combination, which involve a synthesis of distinct sensory cues, cognitive penetration is alleged to involve a synthesis of sensory and *cognitive* information within perceptual processing. In such cases, the operations which determine what we see, hear, and feel, are seen to be directly influenced by what we think, want, and expect. Whether or not cognitive penetration (in this strict sense) actually ever occurs remains a matter of dispute (see Pylyshyn, 1999, Firestone & Scholl, 2016). However, if it does, then such effects may be naturally accommodated by inferentialism in much the same way cases of cue combination are seen to be: perceptual systems form hypotheses about environmental properties under uncertainty, and take on new representational content derived from other sources—in this case *cognitive* sources—to inferentially determine the organism's surroundings. On this picture, cognition 'penetrates' perception (see Hohwy, 2013).

In the case of pain, candidate forms of cognitive penetration mainly concern the influences of expectations on judgements about pain intensity. Such cases are primarily exemplified by placebo and nocebo effects, which are frequently regarded as among the best available evidence for an inferential view of pain (Seymour & Mancini, 2020, Tabor et al., 2017, Shevlin & Friesen, 2021). In what follows, I will focus on the placebo effect in particular. Placebos are a form of dummy treatment, such as a sugar pill or sham surgery, without any identifiable inherent therapeutic properties. Upon administration, however, a placebo may lead patients to report (oftentimes significant) pain relief, exceeding the relief reported by control groups which receive neither active nor placebo treatment (Wager & Fields, 2013). The most central determinants of this effect are generally seen to be (i) conditioning of pain relief with explicit sensory cues, (ii) (conscious) expectations of treatment efficacy, and (iii) the psychosocial context surrounding treatment (Price et al., 2008).

To give a concrete example, illustrating the central role of expectations, consider a type of placebo effect induced in the 'open-hidden paradigm'. In this setting, a subject is either administered treatment in an 'open' condition, where they receive a drug from a healthcare

professional in full view (and who may also verbally assure the subject of the drug's analgesic potency). Or, alternatively, the subject is administered treatment in a 'hidden' condition, where they receive the same drug unawares by means of an automated drug infusion pump (and without any verbal assurance of the drug's analgesic potency). Studies employing the open-hidden paradigm have shown that open administration of a treatment is significantly more effective than hidden administration (*ibid.*). Subjects in the open condition report greater pain relief and require less medication to reach postoperative analgesia than subjects in the hidden condition (Amanzio et al., 2001). The difference in medication needed is seen to reflect the placebo effect. These results suggest that open administration of medical care, and promises of treatment efficacy, make a significant contribution to subjects' treatment response. On a standard interpretation, this is because open treatment creates positive expectations of analgesia, which are linked to pain relief (Wager & Fields, 2013).

Proponents of inferentialism suggest that this 'link' between a cognitive state and pain intensity is inferential: that expectations of pain relief act as premises which inform pain processes in their estimation of imminent threat to the body (Tabor et al., 2017, Seymour & Dolan, 2013). However, this suggestion faces similar difficulties as do candidate cases of cue combination: One difficulty is that it is unclear why we should attribute the placebo effect to a cognitive influence on pain processing and not to a cognitive influence on subjects' judgements of pain (a point which Casser and Clare (2023): 13f. stress at length).⁶ After all, the measurement of a placebo effect usually requires patients to rate their pain or indicate that they have reached a satisfactory analgesic state, which requires them to make a judgement about the intensity of their pain experience. Importantly, however, the judgement process can be influenced by various factors. As is widely recognised, explicit anticipation of pain reduction can lead subjects to (i) establish a lower cognitive anchoring point in their assessment of pain, (ii) overemphasise moments of low pain when judging overall pain experiences, (iii) report what they think the experimenter expects, and (iv) report what they would like to be the case (Wager & Fields, 2013). As such, various researchers have noted that the placebo effect may be explained (at least in part) by differences in reporting decisions which, in turn, reflect differences in cognitive bias, rather than differences in inferential pain analyses.

A second difficulty is that even if the placebo effect did involve cognitive penetration in pain processing, it is unclear why we should think this integration involves inferential, as opposed to associative transitions. As is widely acknowledged, numerous placebo effects are the result of conditioning: a form of associative learning whereby contextual or chemical cues surrounding treatment become associated with pain relief. Some of the clearest evidence for conditioning effects of this kind come from studies involving pharmacological methods. For instance, administering an inert placebo after consecutive administration of an active analgesic, such as morphine, produces a correlated, 'morphine-like' analgesic response (Amanzio & Benedetti, 1999, Colloca & Benedetti, 2006). More generally, conditioning methods of pairing treatment context (e.g. drug, hospital, nurse) with analgesic response can generate future pain relief in virtue of association with the conditioned cue. But if that's true, then a variety of placebo effects can be explained by appealing to associative mechanisms, rendering the introduction of additional, inferential mechanisms in pain processing redundant.

⁶ In a different vein, Klein (2024) has recently argued that many types of cognitive influence on pain processing are best explained in terms of 'transducer calibration', and hence are not even candidates for inferential processing.

As with cases of cue combination, this is not to say that there are no cases of placebo analgesia, or placebo and nocebo effects more broadly, which *could* potentially motivate the thought that pain processing is inferential. However, placebo analgesia as a phenomenon is itself not obviously in need of an inferential interpretation.

3.5. Perceptual Constancy and Invariance

A final pair of possible motivations for inferentialism concerns perceptual constancy and invariance: cases in which our perception of the world remains stable despite significant variation in perceptual conditions. *Perceptual constancies* are mechanisms which ensure uniformity when it comes to a perceptual object's or scene's properties, even when proximal sensory stimulation varies greatly. This allows us, for example, to see a bowl as being 'the same colour' despite significant local differences in illumination across its surface (see Cohen, 2015). To explain how perceptual systems achieve such constancy in the face of variation, it is assumed that they establish the state of the organism's environment inferentially, allowing them to discount variations in their inputs (Rescorla, 2015a).

Invariance, on the other hand, is a feat of object recognition. It characterises our ability to recognise (specific) objects (rather than their properties) across different viewing conditions. This enables us, for example, to recognise something as a fire hydrant, or to recognise a specific person's face, despite the fact that we have never seen either under the particular perceptual conditions we currently find ourselves in. To explain how perceptual systems do this, it is assumed that they infer the identity of a given object in a 'matching process' whereby incoming sensory inputs are compared to stored representations of entities encountered in the past (see Ritchie, 2022).

However, even if it was agreed that perceptual constancy and invariance are straightforwardly indicative of inferential underpinnings, there is seemingly no evidence for analogues of either phenomenon in the case of pain. As evidence for perceptual constancy, what we would expect to find is that our pain experiences represented features of bodily damage as uniform in some way, even when the inputs to the system vary greatly. We would expect to find, for example, that we'd experience pain as 'of the same intensity' or 'at the same location' even when nociceptive information pertaining to the representata fluctuates. However, even if it was granted that pain experiences are representational, there are, to my knowledge, no established effects of this kind (Burge, 2010).

Similarly, as evidence for invariance, what we would expect to find is that we have an ability to 'recognise' specific pains (or specific kinds of pain perhaps) across time and different perceptual conditions. However, insofar as one can make sense of this ability at all, it doesn't seem that we ever recognise or identify pains in anything like the way in which we recognise visual objects. For one thing, we do not pick out pains from other objects in pain experience the way we might pick out a fire hydrant from a trash can in visual experience. After all, all we are ever presented with in pain experience are 'pains', so what role can object recognition even play here? For another, it isn't clear what 'differences in perceptual conditions' would even amount to for pain. We are not separated from our pains the way in which we are separated from the objects of our visual experiences, and hence it is less clear what differences there might be across which we could exercise an ability to recognise pains. Having said that, even if these considerations do not

rule out the possibility of invariance in pain experience, the absence of established candidate cases surely speaks against invariance as an argument in favour of inferentialism about pain.

And so, neither perceptual constancy nor invariance successfully motivate inferential mechanisms in pain processing for the (admittedly) uninteresting reason that there is no evidence for either kind of phenomenon in the case of pain. Unless such evidence can be gathered, the postulation of inferential mechanisms remains explanatorily redundant.

4. TRANSDUCING PAIN

I have argued that standard arguments for inferentialism do not successfully carry over to pain. More specifically, I have argued that (i) pain processing does not solve an underdetermination problem, that (ii) candidate cases of illusion, cue combination, and cognitive penetration fail to motivate a problem of inference for pain, and that (iii) potential further indicators of inferential underpinnings, such as constancy mechanisms and invariance, are absent from pain processing. If this is correct, then pain does not manifest any of the prominent explananda which inferentialism is intended to explain, rendering the introduction of inferential mechanisms gratuitous and, hence, unjustified.

Having said that, the debate over inferentialism remains (to a large extent) subject to empirical fortune, and so proponents of inferentialism about pain are not without options: they may, for example, wish to try and identify more compelling pain-analogues to inferential perceptual processes than those which have been discussed so far. Alternatively, they may wish to show that there are pain phenomena which motivate inferential processing in largely idiosyncratic ways, independent of pain's similarity or dissimilarity to canonical forms of perception. However, in the absence of concrete suggestions, the burden remains on the proponent of inferentialism to show why we need to postulate unconscious inferences in order to account for the operations of (human) pain systems.

And yet, one may rightfully wonder what a non-inferential pain system would look like exactly. In what remains of this paper, I would like to sketch at least one option for how we might think of such a system. A crude way of understanding the difference between an inferential and a non-inferential system is in terms of how 'smart' these systems are. Inferential systems are seen to be 'very smart' (see Fodor, 1985, Gregory, 1970): they actively interpret their inputs on the basis of information they themselves contribute and produce outputs which typically contain more information than the proximal inputs they receive. Perceptual systems are smart, inferential systems *par excellence*. Non-inferential systems, by contrast, are seen to be 'pretty dumb' (Fodor, 1985): instead of interpreting or contributing anything, they merely map inputs to outputs without any complex intermediary steps. There is no more information contained in the outputs they produce than in the inputs they receive. Such 'dumb' systems are sometimes identified with *transducers* (see Pylyshyn, 1984: ch. 6, Klein, 2024).

The notion of a transducer is familiar from biology where it denotes a system which converts one form of energy into another (Loewenstein, 1960). Sensory receptors, such as mechanoreceptors in the skin, for example, transduce (and hence convert) mechanical energy into electrical nerve signals. In this way, transducers re-transmit and convert information, but do so without manipulating (or computing over) the information they are transmitting. Insofar as philosophers of cognitive science regularly regard this process as *the* contrast notion to inferential

(and computational) processing (see Fodor, 1983: 41), it naturally suggests itself as an alternative characterisation of pain processing. On this view, pain's sensory elements are transducers which transmit (and perhaps convert) their inputs, but do so without interpreting or manipulating them.

However, if transduction is indeed the right way to conceive of pain's sensory processing, I think we need to add a couple of qualifications which differentiate pain systems from canonical transducers. The primary reason to think this is that transduction is ordinarily considered to be a *pre*-perceptual process. It is 'the bridge' between the environment and sensory processing proper (Pylyshyn, 1984). However, if my suggested alternative explanations of the thermal grill illusion (Section 3.2.) and phantom limb pain (Section 3.3) are roughly correct, then there is reason to think that pain systems (at least occasionally) take *post*-perceptual inputs. In the case of the TGI, for example, I suggested that the inputs to the pain system are the outputs of a tactile system (i.e. the mistaken temperature estimate of the stimulus). If so, then the transduction process of pain may at times begin *after* the (inferential) process of early tactile processing has been concluded.

A further, independent reason to resist the identification of pain systems with canonical, pre-perceptual transduction is that the relationship between pain states and the activity of candidate (pre-perceptual) sensory receptors is complicated at best. The firing patterns of so-called nociceptors, for example, which are a functionally defined class of receptors sensitive to noxious stimulation, is not straightforwardly related to pain and involves a fair amount of modulation at various stages of spinal processing (Wall, 1979, Heinricher & Fields, 2013). As such, the transduction process cannot be as simple as a mere one-way transmission from neural receptor to pain state. Rather, the transduction process (while non-inferential) must be more complex than what is usually ascribed to 'dumb' transducers.

For these reasons, one would have to conceive of pain's sensory processing as an 'idiosyncratic' form of transduction: not as a bridge between environment and sensory processing, but between perception and cognition. On this view, pain processing often starts where (early) perceptual processing comes to an end:⁷ perceptual systems, such as the tactile system, process and (inferentially) interpret sensory information to determine the state of the organism's (bodily) environment. In some such cases, the perceptual system determines (rightly or wrongly) that the stimulus is exceedingly hot, cold, forceful, acidic, and so forth, in which case its output triggers the activity of a pain system, whose sensory elements transduce the sensory information they receive 'upstream', as it were, where, depending on the influences of various cognitive factors, they ultimately lead to a pain state and the initiation of the organism's protective behaviours. In this way, pain's sensory processing is distinct from that of paradigmatic perceptual modalities and as well as that of canonical transducers.

While this proposal remains tentative and could be fleshed out in a number of different ways, it illustrates that pain is, in many ways, a largely idiosyncratic modality which cannot be understood by mere analogy to paradigmatic perceptual modalities but deserves sustained philosophical attention in its own right. In view of this upshot, I am hopeful that future engagement with these issues will allow for a more detailed account of what pain's mental architecture is like.

⁷ However, pain systems plausibly receive earlier inputs as well, as e.g. in situations where the heat pain of touching the hot stove almost reflexively causes limb withdrawal.

ACKNOWLEDGEMENTS

Earlier versions of this paper have been presented at the 2023 Central APA in Denver, CO, the SSPP's 113th annual meeting in Mobile, AL, and at the 2022 graduate philosophy conference in Austin, TX. I am grateful to all audience members for their helpful questions and comments. I am especially indebted to Gabe Siegel, Nada Gligorov, and Matthew Fulkerson for their insightful commentaries at these occasions. I would also like to thank Luca Barlassina, Ray Buchanan, Anabel Casser, Josh Dever, Andrea Rivadulla Duró, Ben Henke, Hans Kamp, Alex Kerr, Colin Klein, Kevin Lande, Michelle Montague, Ian Phillips, Anne Quaranto, Mark Sainsbury, Matt Vermaire, and an anonymous reviewer of this journal for detailed feedback on earlier drafts. Special thanks to Sam Clarke for his unwavering support of this project despite the fact that it diverges significantly from views we jointly expressed in the past.

REFERENCES

1. Amanzio, M., & Benedetti, F. (1999). Neuropharmacological dissection of placebo analgesia: expectation-activated opioid systems versus conditioning-activated specific subsystems. *Journal of Neuroscience*, *19*(1), 484-494.
<https://doi.org/10.1523/JNEUROSCI.19-01-00484.1999>
2. Amanzio, M., Pollo, A., Maggi, G., & Benedetti, F. (2001). Response variability to analgesics: a role for non-specific activation of endogenous opioids. *Pain*, *90*(3), 205-215. DOI: 10.1016/S0304-3959(00)00486-3
3. Anchisi, D., & Zanon, M. (2015). A Bayesian perspective on sensory and cognitive integration in pain perception and placebo analgesia. *PLoS one*, *10*(2), e0117270.
<https://doi.org/10.1371/journal.pone.0117270>
4. Armstrong, D.M. (1968). *A Materialist Theory of the Mind*. London: Routledge & Kegan Paul.
5. Bach, P., Becker, S., Kleinböhl, D., & Hölzl, R. (2011). The thermal grill illusion and what is painful about it. *Neuroscience letters*, *505*(1), 31-35.
<https://doi.org/10.1016/j.neulet.2011.09.061>
6. Barlassina, L., & Hayward, M. K. (2019). More of me! Less of me!: Reflexive imperativism about affective phenomenal character. *Mind*, *128*(512), 1013-1044.
<https://doi.org/10.1093/mind/fzz035>
7. Baysan, U. & Macpherson, F. (2017), 'Aristotle's Illusion'. In F. Macpherson (ed.), *The Illusions Index*. Retrieved from <https://www.illusionsindex.org/i/aristotle>.
8. Block, N. (2014). Seeing-as in the light of vision science. *Philosophy and Phenomenological Research*, *89*(3), 560-572. <https://doi.org/10.1111/phpr.12135>
9. Block, N. (2022). *The Border Between Seeing and Thinking*. Oxford: Oxford University Press.
10. Bouhassira, D., Kern, D., Rouaud, J., Pelle-Lancien, E., & Morain, F. (2005). Investigation of the paradoxical painful sensation ('illusion of pain') produced by a thermal grill. *Pain*, *114*(1-2), 160-167. doi: 10.1016/j.pain.2004.12.014
11. Brodie, E. E., Whyte, A., & Niven, C. A. (2007). Analgesia through the looking-glass? A randomized controlled trial investigating the effect of viewing a 'virtual' limb upon phantom limb pain, sensation and movement. *European journal of pain*, *11*(4), 428-436.
<https://doi.org/10.1016/j.ejpain.2006.06.002>

12. Brown, C. A., Seymour, B., Boyle, Y., El-Deredy, W., & Jones, A. K. (2008). Modulation of pain ratings by expectation and uncertainty: Behavioral characteristics and anticipatory neural correlates. *Pain, 135*(3), 240-250. doi: 10.1016/j.pain.2007.05.022
13. Brown, C. A., Seymour, B., El-Deredy, W., & Jones, A. K. (2008). Confidence in beliefs about pain predicts expectancy effects on pain perception and anticipatory processing in right anterior insula. *Pain, 139*(2), 324-332. doi 10.1016/j.pain.2008.04.028
14. Burge, T. (2010). *Origins of Objectivity*. Oxford: Oxford University Press.
15. Casser, L. C. (2021). The function of pain. *Australasian Journal of Philosophy, 99*(2), 364-378. <https://doi.org/10.1080/00048402.2020.1735459>
16. Casser, L., & Clarke, S. (2023). Is pain modular?. *Mind & Language, 38*(3), 828-846. <https://doi.org/10.1111/mila.12430>
17. Cohen, J. (2015). 'Perceptual constancy'. In Mohan Matthen (Ed.) *The Oxford handbook of philosophy of perception* (pp.621-639). Oxford: Oxford University Press.
18. Colloca, L., & Benedetti, F. (2006). How prior experience shapes placebo analgesia. *Pain, 124*(1-2), 126-133. doi: 10.1016/j.pain.2006.04.005
19. Coninx, S. (2021). Strong representationalism and bodily sensations: Reliable causal covariance and biological function. *Philosophical Psychology, 34*(2), 210-232. <https://doi.org/10.1080/09515089.2020.1858476>
20. Craig, A. D., & Bushnell, M. C. (1994). The thermal grill illusion: unmasking the burn of cold pain. *Science, 265*(5169), 252-255. doi: 10.1126/science.8023144
21. Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature, 415*(6870), 429-433. <https://doi.org/10.1038/415429a>
22. Fardo, F., Beck, B., Allen, M., & Finnerup, N. B. (2020). Beyond labeled lines: A population coding account of the thermal grill illusion. *Neuroscience & Biobehavioral Reviews, 108*, 472-479. doi: 10.1016/j.neubiorev.2019.11.017
23. Finlay, B. L. (2019). The neuroscience of vision and pain: evolution of two disciplines. *Philosophical Transactions of the Royal Society B, 374*(1785), 20190292. doi: 10.1098/rstb.2019.0292
24. Firestone, C., & Scholl, B. J. (2016). Cognition does not affect perception: Evaluating the evidence for "top-down" effects. *Behavioral and brain sciences, 39*, e229. doi: 10.1017/S0140525X15000965
25. Fodor, J. A. (1983). *The modularity of mind*. MIT Press.
26. Fodor, J. A. (1985). Precis of the modularity of mind. *Behavioral and brain sciences, 8*(1), 1-5. doi: 10.1017/S0140525X0001921X
27. Fodor, J. A., & Pylyshyn, Z. W. (1981). How direct is visual perception?: some reflections on Gibson's " Ecological Approach". *Selected Readings in the Philosophy of Perception, 167*. [https://doi.org/10.1016/0010-0277\(81\)90009-3](https://doi.org/10.1016/0010-0277(81)90009-3)
28. Gibson, J. (1979). *The Ecological Approach to Visual Perception*. Boston, Houghton Mifflin.
29. Gligorov, N. (2017). Don't worry, this will only hurt a bit: The role of expectation and attention in pain intensity. *The Monist, 100*(4), 501-513. <https://doi.org/10.1093/monist/onx025>
30. Green, B. G. (2002). Synthetic heat at mild temperatures. *Somatosensory & motor research, 19*(2), 130-138. <https://doi.org/10.1080/08990220220220131524>

31. Green, B. G. (1977). Localization of thermal sensation: An illusion and synthetic heat. *Perception & Psychophysics*, 22(4), 331-337. <https://doi.org/10.3758/BF03199698>
32. Gregory, R. L. (1970). *The intelligent eye*. London: Weidenfeld & Nicolson.
33. Hatfield, G. (2002). 'Perception as Unconscious Inference'. In Dieter Heyer & Rainer Mausfeld (Eds.) *Perception and the Physical World* (pp.113-143). New York: John Wiley & Sons.
34. Heinricher, M.M. and H.L. Fields (2013). Central Nervous System Mechanisms of Pain Modulation. In S. McMahon, M. Koltzenburg, I. Tracey, & D. Turk (Eds.) *Wall and Melzack's Textbook of Pain* (129-142). Philadelphia: Elsevier Saunders.
35. Von Helmholtz, H. (1867). *Handbuch der Physiologischen Optik*. Leipzig: Voss.
36. Hohwy, J. (2013). *The predictive mind*. Oxford: Oxford University Press.
37. Hsiao, S. (2008). Central mechanisms of tactile shape perception. *Current opinion in neurobiology*, 18(4), 418-424. doi: 10.1016/j.conb.2008.09.001
38. Jacobson, H. (2017). 'Pain and cognitive penetrability'. In Jennifer Corns (Ed.) *The Routledge Handbook of Philosophy of Pain* (pp. 266-275). Abingdon: Routledge. <https://doi.org/10.4324/9781315742205>
39. Jenkin, Z. (2022). Crossmodal basing. *Mind*, 131(524), 1163-1194. <https://doi.org/10.1093/mind/fzac030>
40. Klein, C. (2007). An imperative theory of pain. *The Journal of Philosophy*, 104(10), 517-532. <https://doi.org/10.5840/jphil2007104104>
41. Klein, C. (2015). *What the Body Commands: The Imperative Theory of Pain*. MIT Press. doi: <https://doi.org/10.7551/mitpress/10480.001.0001>
42. Klein, C. (2024). Transduction, calibration, and the penetrability of pain. *Ergo an Open Access Journal of Philosophy*, 10. doi: <https://doi.org/10.3998/ergo.5187>
43. Lederman, S. J., & Jones, L. A. (2011). Tactile and haptic illusions. *IEEE Transactions on Haptics*, 4(4), 273-294. doi: 10.1109/TOH.2011.2
44. Loewenstein, W. R. (1960). Biological transducers. *Scientific American*, 203(2), 98-111. doi:10.1038/scientificamerican0860-98
45. Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: W.H. Freeman.
46. Martínez, M. (2011). Imperative content and the painfulness of pain. *Phenomenology and the Cognitive Sciences*, 10, 67-90. doi: 10.1007/s11097-010-9172-0
47. Melzack, R., Wall, P. D., & Ty, T. C. (1982). Acute pain in an emergency clinic: latency of onset and descriptor patterns related to different injuries. *Pain*, 14(1), 33-43. doi: 10.1016/0304-3959(82)90078-1
48. Melzack, R. and P.D. Wall 2008 (1982). *The Challenge of Pain*, New York: Penguin Science.
49. Moseley, G. L., & Arntz, A. (2007). The context of a noxious stimulus affects the pain it evokes. *Pain®*, 133(1-3), 64-71. doi: 10.1016/j.pain.2007.03.002
50. Moutoussis, M., Fearon, P., El-Deredy, W., Dolan, R. J., & Friston, K. J. (2014). Bayesian inferences about the self (and others): A review. *Consciousness and cognition*, 25, 67-76. doi: 10.1016/j.concog.2014.01.009
51. Nikolajsen, L. (2013). 'Phantom Limb'. In S. McMahon, M. Koltzenburg, I. Tracey, & D. Turk (Eds.) *Wall and Melzack's Textbook of Pain* (248-255). Philadelphia: Elsevier Saunders.
52. Ongaro, G., & Kaptchuk, T. J. (2019). Symptom perception, placebo effects, and the Bayesian brain. *Pain*, 160(1), 1-4. doi: 10.1097/j.pain.0000000000001367

53. Orlandi, N. (2014). *The innocent eye: Why vision is not a cognitive process*. Oxford: Oxford University Press.
54. Palmer, S. E. (1999). *Vision science: Photons to phenomenology*. Cambridge, MA: MIT Press.
55. Phillips, I. & Firestone (2023). Visual adaptation and the purpose of perception, *Analysis* 83(3): 555-575. <https://doi.org/10.1093/analys/anac060>
56. Price, D. D., Finniss, D. G., & Benedetti, F. (2008). A comprehensive review of the placebo effect: recent advances and current thought. *Annu. Rev. Psychol.*, 59(1), 565-590. doi: 10.1146/annurev.psych.59.113006.095941
57. Pylyshyn, Z. (1984). *Computation and cognition: Toward a foundation for cognitive science*. Cambridge, MA: MIT Press.
58. Pylyshyn, Z. (1999). Is vision continuous with cognition?: The case for cognitive impenetrability of visual perception. *Behavioral and brain sciences*, 22(3), 341-365. doi: 10.1017/s0140525x99002022
59. Ramachandran, V. S., & Altschuler, E. L. (2009). The use of visual feedback, in particular mirror visual feedback, in restoring brain function. *Brain*, 132(7), 1693-1710. doi: 10.1093/brain/awp135
60. Ramachandran, V. S., Rogers-Ramachandran, D., & Cobb, S. (1995). Touching the phantom limb. *Nature*, 377(6549), 489-490. doi: [10.1038/377489a0](https://doi.org/10.1038/377489a0)
61. Ramachandran, V. S. (1988). Perceiving shape from shading. *Scientific American*, 259(2), 76-83. doi: 10.1038/scientificamerican0888-76
62. Rescorla, M. (2015a). 'Bayesian Perceptual Psychology'. In Mohan Matthen (Ed.) *The Oxford Handbook of the Philosophy of Perception*. Oxford: Oxford University Press: 694-716.
63. Rescorla, M. (2015b). The Innocent Eye: Why Vision is Not a Cognitive Process. *Notre Dame Philosophical Reviews*.
64. Rescorla, M. (2020). Perceptual co-reference. *Review of Philosophy and Psychology*, 11(3), 569-589. <https://doi.org/10.1007/s13164-018-0411-6>
65. Ringkamp, M., Srinivasa, R. N., Campbell, J. N., & Meyer, R. A. (2013). Peripheral mechanisms of cutaneous nociception. In S. McMahon, M. Koltzenburg, I. Tracey, & D. Turk (Eds.), *Wall and Melzack's textbook of pain* (pp. 1–30). Philadelphia: Elsevier Saunders.
66. Brendan Ritchie, J. (2022). Recognizing why vision is inferential. *Synthese*, 200(1), 25. doi: 10.1007/s11229-022-03508-1
67. Rock, I. (1975). *An introduction to perception*. New York: Macmillan.
68. Rock, I. (1983). *The Logic of Perception*. Cambridge, MA: MIT Press.
69. Rosenqvist, T. (forthcoming). 'The Pain System is not a Bodily Disturbance Detector'. In Cuevas-Badallo, A., Martín-Villuendas, M., Gefaell, J. (Eds.), *Life and Mind: Theoretical and Applied Issues in Contemporary Philosophy of Biology and Cognitive Sciences*. Springer.
70. Seymour, B. and Dolan, R.J. (2013), 'Emotion, Motivation, and Pain'. In S. McMahon, M. Koltzenburg, I. Tracey, & D. Turk (Eds.), *Wall and Melzack's textbook of pain* (p. 248-255). Philadelphia: Elsevier Saunders.
71. Seymour, B., & Mancini, F. (2020). Hierarchical models of pain: Inference, information-seeking, and adaptive control. *NeuroImage*, 222, 117212. <https://doi.org/10.1016/j.neuroimage.2020.117212>
72. Shevlin, H., & Friesen, P. (2021). Pain, placebo, and cognitive penetration. *Mind & Language*, 36(5), 771-791. <https://doi.org/10.1111/mila.12292>

73. Skrzypulec, B. (forthcoming). 'Pain: Modularity and Cognitive Constitution', *The British Journal for the Philosophy of Science*. <https://doi.org/10.1086/727001>
74. Tabor, A., Thacker, M. A., Moseley, G. L., & Körding, K. P. (2017). Pain: a statistical account. *PLoS computational biology*, *13*(1), e1005142. <https://doi.org/10.1371/journal.pcbi.1005142>
75. Thunberg, T., (1896). 'Förnimmelserna vid till samma ställe lokaliserad, samtidigt pågående köld-och värmeretning'. Uppsala Läkfören Förh.
76. Trommershäuser, J., Körding, K., & Landy, M. S. (2011). *Sensory cue integration*. Oxford: Oxford University Press.
77. Tye, M. (1995a). *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*, Cambridge, MA: The MIT Press.
78. Tye, M. (1995). A representational theory of pains and their phenomenal character. *Philosophical perspectives*, *9*, 223-239.
79. von Helmholtz, H. (1962) [1867]. *Treatise on Physiological Optics*, trans. J. Southall. New York: Dover.
80. Wager, T. D. (2005). Expectations and anxiety as mediators of placebo effects in pain. *Pain*, *115*(3), 225-226. doi: 10.1016/j.pain.2005.03.018
81. Wager, T. D., & Fields, H. L. (2013). 'Placebo analgesia'. In S. McMahan, M. Koltzenburg, I. Tracey, & D. Turk (Eds.), *Wall and Melzack's textbook of pain* (pp. 362–373). Philadelphia: Elsevier Saunders.
82. Wall, P.D. (1979). On the Relation of Injury to Pain. *Pain* *6*(3): 253–64. doi: 10.1016/0304-3959(79)90047-2
83. Wall, P.D. & McMahan, S.B. (1986). 'The relationship of perceived pain to afferent nerve impulses', *Trend in Neurosciences* *9*(C), 254-255. [https://doi.org/10.1016/0166-2236\(86\)90070-6](https://doi.org/10.1016/0166-2236(86)90070-6)
84. Watson, A., El-Deredy, W., Bentley, D. E., Vogt, B. A., & Jones, A. K. P. (2006). Categories of placebo response in the absence of site-specific expectation of analgesia. *Pain*, *126*(1), 115-122. doi: 10.1016/j.pain.2006.06.021
85. Wiech, K., Vandekerckhove, J., Zaman, J., Tuerlinckx, F., Vlaeyen, J. W., & Tracey, I. (2014). Influence of prior information on pain involves biased perceptual decision-making. *Current Biology*, *24*(15), R679-R681. doi: 10.1016/j.cub.2014.06.022
86. Wiech, K. (2016). Deconstructing the sensation of pain: The influence of cognitive processes on pain perception. *Science*, *354*(6312), 584-587. doi: 10.1126/science.aaf8934
87. Yoshida, W., Seymour, B., Koltzenburg, M., & Dolan, R. J. (2013). Uncertainty increases pain: evidence for a novel mechanism of pain modulation involving the periaqueductal gray. *Journal of Neuroscience*, *33*(13), 5638-5646. <https://doi.org/10.1523/JNEUROSCI.4984-12.2013>