

TRANSDUCTION, CALIBRATION, AND THE PENETRABILITY OF PAIN

COLIN KLEIN

The Australian National University

Pains are subject to obvious, well-documented, and striking top-down influences. This is in stark contrast to visual perception, where the debate over cognitive penetrability tends to revolve around fairly subtle experimental effects. Several authors have recently taken up the question of whether top-down effects on pain count as cognitive penetrability, and what that might show us about traditional debates. I review some of the known mechanisms for top-down modulation of pain, and suggest that it reveals an issue with a relatively neglected part of the cognitive penetrability literature. Much of the debate inherits Pylyshyn's stark contrast between transducers and cognition proper. His distinction grew out of his running fight with Gibson, and is far too strong to be defensible. I suggest that we might therefore view top-down influences on pain as a species of transducer calibration. This provides a novel but principled way to distinguish between several varieties of top-down effect according to their architectural features.

1. Top-Down Effects on Pain

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. The nail was then pulled out from below. 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, Hassan, & O'Connor 1995)

Contact: Colin Klein <colin.klein@anu.edu.au>

This vignette presents us with a mystery. Nobody doubts that the poor builder was in severe pain. Yet while his reputation on the job site was irreparably damaged, *physically* speaking he was completely unharmed.

How does this sort of thing happen? Here is a tempting story—one that I think is broadly correct, though the details will come to matter quite a bit. Having jumped down, our builder saw the nail sticking through his boot. (As the picture accompanying the case report makes clear, the sight would have been a gruesome one.) He thus formed the belief that he was grievously injured, despite the lack of substantial nociceptive input from his foot. That belief was a purely cognitive state. Yet it caused him to have certain perceptual experiences commensurate with it: pain followed. Call this a ‘top-down’ explanation of the builder’s pain: it posits that states like beliefs (that we normally think of as relatively high up in some hierarchy of cognitive states) end up affecting perceptual states (that we normally think of as lower down in the same hierarchy).

The reason why this explanation is so natural is that top-down effects on pain are so common. Both Gligorov (2017) and Shevlin and Friesen (2021) have recently called attention to placebo analgesia as a putative case of top-down influence.¹ Tell someone that the inert pill they are about to receive will diminish their pain, and often enough it will do just that. Conversely, emphasizing that a procedure is likely to hurt tends to make it *more* painful. Both effects are evoked so reliably that they affect clinical practice: placebo effects are *so* easy to evoke that placebo control has become a standard against which claims of drug effectiveness must be measured, while nocebo effects raise serious practical issues for informed consent.²

Hypnosis similarly has a notable (if variable) influence on pain perception (Patterson 2004; Patterson & Jensen 2003). Importantly, hypnosis has been used to ameliorate the pain of labor, burn debridement, and bone marrow aspiration—all intense pains that are often resistant to other interventions. Yet hypnosis is, on the face of it, just talk. The diminution of pain can be more dramatic still. It has been long known that many severe injuries are initially painless, often for a period of hours. Wall argues (conclusively in my opinion) that this is an adaptive response: that the role of pain is to limit motion, and the downward suppression is elicited by appraisal that the circumstances are such that “treatment of the injury does not have the highest biological priority” (Wall 1979b:

1. Philosophical treatments of pain have often said something about Placebo, so this is far from exhaustive. Hardcastle (1997) also cites many useful cases; see especially 392ff.

2. Benedetti et al. (2007) reviews literature on both the placebo and nocebo effects. For useful discussions of ethical issues surrounding informed consent and nocebo effects, see Wells and Kaptchuk (2012) and Gligorov (2018).

298). Belief about the need to escape to safety can thus also modulate the pain of severe injury (Wall 1979a; 1979b; 2000).³

In addition to these dramatic effects, there is considerable evidence for more complex modulation of pain. Work on the “biopsychosocial” (Gatchel, Peng, Peters, Fuchs, & Turk 2007) or “biopsychomotor” (Sullivan 2008) models of pain suggests that pain is influenced not only by physical processes but also by social and psychological mediators—cognitive processes *par excellence*. Felt control Staub, Tursky, and Schwartz (1971) and self-efficacy (Litt 1988) over painful stimuli is known to diminish the intensity of pain, and this appears to be modulated by a suite of cognitive factors (Thompson 1981). Physiotherapists have emphasized the role of false beliefs about the origins of pain in the persistence of chronic pain, and hence the role of patient education in diminishing pain (Butler & Moseley 2013; Moseley 2002; Moseley & Butler 2015). Social and contextual factors also play an important role in moderating placebo analgesia (Atlas 2021).

Hence Shevlin and Friesen conclude that “. . . it appears that one’s beliefs and expectations with regards to the treatment one is receiving, as well as one’s awareness of the administration, have a significant impact on one’s immediate experience of pain relief” (2021: 8). What do these top-down influences show. Shevlin and Friesen (2021) have suggested that effects like these are evidence for something like *cognitive penetrability* of pain.

Cognitive penetrability is philosophically contentious notion. Much of what follows will be concerned with exploring different definitions. But we need a place to start, and I think the definition given by Stokes (2013) captures a common low-commitment starting point:

A perceptual experience *E* is cognitively penetrated if and only if (1) *E* is causally dependent upon some cognitive state *C* and (2) the causal link

3. Melzack, Wall, and Ty (1982) found that 37% of emergency room patients had some substantial painless period. Anecdotally, every doctor I’ve asked assures me that it is so common as to make individual cases unremarkable. The most well-known discussion of painless injury is likely Beecher (1956)’s discussion of the injured at Anzio, but I follow Wall in doubting Beecher’s interpretation of the mechanisms.

It is important to note that this is not a matter of minor injuries overlooked in the heat of the moment. Melzack et al., speaking of emergency patients they studied, note that

They took appropriate steps to go to the hospital, and were not confused or in shock at any time. They were fully aware of the extent of their injuries, and therefore astonished at the lack of pain until it finally set in. The patients who arrived at the hospital without pain and remained pain-free repeatedly expressed their surprise that a major, obvious injury such as a severe laceration would cause no pain. (1982: 41)

In other words, these patients in are fully aware of their injuries and attend to them, yet do not feel pain.

between *E* and *C* is internal and mental. (Stokes 2013: 650; compare also Stokes 2012: 479)

Clause 1 is straightforward, as is the first conjunct of clause 2. If I desire a drink of water and turn my head to look for the glass, that changes my perceptual experience, but not in a way anyone worries about. One may complain about the boundaries of ‘internal’—if I think about magpie attacks and get scared and then get indigestion, that’s all within the boundaries of the skin—but leave it for now.

The second conjunct of clause 2 of is less clear, and hence where most of the action has been. Since Pylyshyn (1984), one of the key escape clauses for putative cases of penetration has been the effect of *attention*. Merely attending to one stimulus rather than another might change how well one of them is processed, or even if one raises to the level of consciousness at all. The idea is usually that this is merely an inner analogue of turning one’s head. Penetration demands something more: that the effect on perception is in virtue of, and specific to, the *content* of our beliefs and desires.

Now, most of the debate around cognitive penetrability has been around the penetration of visual perception. This has been a fraught debate. For what it’s worth, I am convinced by recent critiques arguing that evidence in favor of penetration is of relatively poor quality (Firestone & Scholl 2016; Machery 2015). Clause 2 of Stokes’s definition gives a lot of wiggle room for those who would be skeptical: attentional effects are easy to find and easier to posit. Furthermore, as Firestone and Scholl (2016) point out, the more plausible cases of visual penetrability are (by and large) fairly subtle and fairly rare. Much of the debate thus ends up focusing on the details of a handful of cases.

As Shevlin and Friesen point out, pain is different: as detailed above, pretty much *everybody* accepts that there are top-down influences, so the only question is whether they fit the definition. Shevlin and Friesen also note that the case of pain (and bodily sensations more generally) might not generalize to sensations like vision, because there is “a fundamental functional and even architectural distinction to be drawn between these two families of mental states, perhaps reflecting distinct evolutionary histories” (2021: 14). So pain is its own case, and it seems that we do have good *prima facie* evidence that there are top-down effects that don’t just amount to (e.g.) someone paying more or less attention to the pain they’re in.

The case of pain is not entirely clear-cut, however. There is quite a lot about pain perception that doesn’t fit the penetrability story. So, for example, a great deal of chronic pain persists despite the sincere and accurate beliefs of those affected that there is no tissue damage (or any other issue outside that of the pain system). Similarly, just having the powerful *desire* that you not be in pain does not seem to affect pain in the slightest. Consider, by way of concrete example,

phantom limb pain. Patients know they do not have an arm, yet that belief seems powerless to stop the feeling of pain. This seems to be precisely analogous to visual illusions: knowing that the two lines in the Müller-Lyer illusion are in fact the same length doesn't keep them from looking the way they do. This would suggest that we ought to draw the same conclusion from phantom limb pain: pain processing is in fact modular and *impenetrable* due to its encapsulation.⁴

So it seems like there is some real debate about whether the top-down influences seen on pain really count as cognitive penetration in the full sense of the term. The present paper will argue that this debate is live because the broad, interesting top-down influences on pain actually belong to another phenomenon altogether, which I will call *transducer calibration*. Like cognitive penetration, transducer calibration allows for complex effects on perceptual processes that are ultimately driven by a variety of different cognitive states—but these effects do not (I claim) threaten the traditional architectural distinctions that penetrability is meant to threaten. For, as opponents of penetrability argue, the effects are on transducers, which are traditionally outside the cognitive system—*except* that appeal to transducers is usually meant to end the argument, whereas I will show that transducer calibration is complex, potentially widespread, and far more interesting than opponents of cognitive penetration have typically credited it with being.

The argument is thus dialectically complex, and will take a bit of time to unfold. Before I begin, two important caveats. First, the goal of the paper is not really to adjudicate between different debates about what cognitive penetrability *is*; I need enough to show that transducer calibration *isn't* penetration, and that should be true on a variety of accounts. One of the major advances of this literature in the past 40 years has been to disentangle various criteria that were jointly proposed for penetration. One of the minor aims of the paper is to show that this disentanglement opens the way for the identification of other, equally interesting, phenomena.

Second, I do not intend my account to cover all putative top-down effects on pain. So, for example, in an elegant series of studies Wiech et al. (2014) show that base rate information affects the judgment of shocks as of low or high intensity. Both Shevlin and Friesen (2021) and Casser and Clarke (2022) both discuss this, and come down on different sides. The dynamic replicates that of traditional debates about vision: one can read Wiech et al. as showing either penetration of an early perceptual process *or* the effect of knowledge on postperceptual deci-

4. See also Casser and Clarke (2022), who argue that the modularity of pain processing ought to be a 'default assumption'. They suggest that this would either make pain impenetrable (if there is a single module) or else allow for penetrability "provided that this penetration simply occurs at the joints between independently posited systems, influencing the outputs of lower-level modules before these are taken as input by higher-level systems" (Casser & Clarke 2022: 11).

sion-making. Similarly, everyone agrees that there are complex and widespread influences of emotion, anticipation, and personality on reactions to pain; some of the effects of placebo analgesia (for example) can certainly be attributed to those (Wiech, Ploner, & Tracey 2008). The claim of the paper is rather than insofar as there are many of these large, striking effects on pain, they belong to the phenomenon of transducer calibration.

My argument will proceed as follows. I begin by arguing that the downward effects on pain for which there are good evidence are downward effects on the *transducers* of pain, rather than on any modular processing of pain. This will come in two steps. In the conceptual step (Section 2), I distinguish transducers from early perception proper; this will also flesh out what's normally at stake in discussions of cognitive penetrability. The empirical step (Section 3) reviews the pathways for downward modulation, and argue that these fit best with a story on which transducers are modulated. The idea that there might be nontrivial modulation of transducers is neither familiar nor obvious; in the course of defending against some objections (Section 4), I'll flesh out why transducer calibration is a distinct and interesting class of downward influence. Finally (Section 5) I'll step back to suggest that transducer calibration may be more widespread than one might suppose, and suggest that this might have several interesting consequences for thinking about cognitive architecture more broadly.

2. Cognitive Penetration and Cognitive Architecture

The broad idea that cognition might influence perception is an old one and appears in many guises (Stokes 2013). The current debate—and, arguably, widespread use of the term itself—arises from Pylyshyn (1984)'s defense of a computational theory of cognition. In what follows, I will lean liberally on Pylyshyn's (1984) and (1999), both of which present the idea of cognitive penetrability in the course of arguing against it.

The architectural picture in the background of Pylyshyn's arguments has been remarkably influential. It is closely related to Fodor (1983)'s defense of a modular theory of mind, which has been arguably even more influential within philosophy. Note that everything from this point out will primarily be of interest to those who adopt a broadly computationalist theory about the mind. Cognitive penetrability sometimes comes up in the course of arguing *against* computationalism, by showing that the architectural boundaries drawn below are illegitimate. I'm not doing that here. If that is your goal, you should probably look to someone other than Pylyshyn for an architectural story within which to frame the debate.

Suppose I gaze upon a panda, recognize it as such, and reflect that if pandas weren't so cute, they would have gone extinct long ago. Pylyshyn identifies three important stages in this process.

First, there is an initial step of *transduction*, in which light from the panda falling on the retina is transformed into a suitable symbolic representation. This is a single step from the computational point of view: a special sort of primitive operation that connects inside to outside.

Next, there are multiple *symbolic* steps. In these, computational processes—sensitive, by definition, only to the formal/syntactic/computational properties of symbols rather than their contents—take symbols in, manipulate them, and pass the resulting symbols on to other processes. Much of cognitive science is concerned with these steps: the particular ways in which edges and other features are extracted, combined with stored information, and used to categorize objects. Finally, there are the *intentional* steps: ones that involve beliefs about the cuteness and haplessness of pandas, and that are combined with other beliefs in an inferential, truth-preserving way.

Each step, says Pylyshyn, requires a different sort of explanation. The explanation of transducers is an engineering problem. As Pylyshyn puts it, "Like all primitive operations of the functional architecture, the transducer fundamentally is a physical process: that is, its behavior is explainable...in terms of its intrinsic properties—physical, chemical, biological, and so on" (Pylyshyn 1984: 148). Explaining transducers is not part of cognitive science proper, he says. Rather, cognitive scientists *assume* appropriate transducers. The explanatory task for cognitive science starts when the world is translated into symbols and ends when symbols get transduced back into action. Similarly, intentional transitions can be explained by the meanings and truth-conditions of the contents involved. Here we can rely on logic and personal-level psychology to explain success and failure.

Hence cognitive science (says Pylyshyn) is primarily concerned with the symbolic, computational step in the middle. It is here that all of the classic work gets done—the determination of algorithms, the delineation of representational formats and the postulation of computational primitives and modules. The fundamental distinction between the symbolic and the inferential, Pylyshyn insists, lies in the nature of the state-transitions. Computational explanations take symbols and transform them according to rules defined only on their formal properties (however that is understood). This is not constrained by, and often fails to follow, the rational norms of inference that govern the intentional step. Indeed, Pylyshyn (1999) often goes to lengths to emphasize ways that early vision *fails* to make rational inferences. Discussing the Kanizsa amodal completion figure, for example, Pylyshyn notes that the seen completion is neither the simplest nor the most likely. Hence whatever is going on in early

vision “follows complex principles of its own—that are generally not rational principles, such as semantic coherence or even something like maximum likelihood” (1999: 345).

One of the important ways in which the symbolic is not rational is that it cannot take into account all available information. Rational inference is *isotropic*: it can (in principle) take into account any belief whatsoever, so long as that belief bears a semantic relationship to the topic at hand. As Fodor puts it:

By saying that confirmation is isotropic, I mean that the facts relevant to the confirmation of a scientific hypothesis may be drawn from anywhere in the field of previously established empirical. . . truths. Crudely: everything that the scientist knows is, in principle, relevant to determining what else he ought to believe. In principle, our botany constrains our astronomy, if only we could think of ways to make them connect. (Fodor 1983: 105)

Conversely, early vision appears to be strongly *anisotropic*: even things you know very well don’t seem to influence vision. We know that the two lines of the Müller-Lyer illusion are the same length, yet we cannot bring this belief to bear. The lines still look unequal: a vivid demonstration of the the anisotropy of early vision.

Finally, the distinction between the symbolic and the intentional also explains why Pylyshyn is so concerned with cognitive penetrability. Pylyshyn takes the “essence” of cognitive penetration to be “an influence that is coherent or quasi-rational when the meaning of the representation is taken into account” (Pylyshyn 1999: 365). Yet Pylyshyn’s story relies on their being *some* orderly division between the merely symbolic and the intentional.⁵ Cognitive penetrability happens when some belief influences a computational, symbolic bit—that normally only cares about symbols—in virtue of its semantic content.

As he puts it “The criterion proposed in this section is a direct consequence of a view that might be called *the basic assumption of cognitive science*. . . that there are at least three, distinct, independent levels at which we can find explanatory principles in cognitive psychology” (Pylyshyn 1984: 133). Or, put another way, early vision is just defined as the anisotropic bits. If they can be affected by belief,

5. This can’t be a strict division: at some point in the chain, transitions according to symbolic properties must *also* be truth-preserving. That is how the intentional can be explained by the computational. The point is that talk of early vision only makes sense if there are some transitions that don’t work like this.

then they are isotropic after all. After all, belief-formation is isotropic, so once you get one belief affecting early vision you get them all. The jig is up.⁶

In an excellent discussion, Burnston and Cohen (2015) argue that debates about modularity and cognitive penetrability have tended to move away from the broad idea anisotropy to a more restrictive conception of modularity—one on which, roughly speaking, modules only have access to information from earlier processing areas. So, for example, Deroy (2013: 93) says that cognitive impenetrability holds just in case “. . . perception depends only on its own rules of processing rules [sic] and on the kind of input it receives, and is independent from higher cognitive contents.” Burnston and Cohen (2015) point out that this restrictive requirement is precisely what generates a lot of the friction around particular cases, particularly any that involve what they call ‘integrative’ cross-modal processes.⁷

Here I closely follow Burnston and Cohen (2015)’s excellent discussion of the debate. They claim—and I find this very plausible—that all putative instances of visual penetration involve influence from information that is semantically coherent but still anisotropic. One of the standard examples of cognitive penetrability is meant to be effects of the typical color of an object on perceived color of instances: that cutouts of a heart are seen as more red than other neutral shapes, or that a picture of a neutral-gray banana will be seen as slightly yellow (Delk & Fillenbaum 1965). Yet the perceived color of bananas is affected by knowledge about the characteristic color of bananas, that is still only one piece of information in the whole doxastic system that is relevant. Note that this feature is, in some sense, independent of whether you think that these experimental effects actually occur: the point is rather than even the most plausible candidates for penetration are still anisotropic.

Even if the influence is semantically coherent, then, the threatened collapse of the distinction between central and modular systems does not occur. Limited, isotropic cognitive influence is thus compatible with Pylyshyn’s project of distinguishing isotropic central inference from anisotropic symbolic processing (Pylyshyn 1984: 330). Conversely, the idea of fully isotropic cognitive influence is hard to understand structurally because, as Burnston and Cohen (2015) note, it is hard to see how representations in one format could be arbitrarily translated to influence representations in different, lower-level formats.

6. Note that Pylyshyn’s approach changes over time. Pylyshyn (1984) takes it as a basic methodological principle that there need to be *some* parts of early vision that are merely symbolic. If there’s not, then the whole project falls apart. By the time of his (1999), Pylyshyn softens a bit and appears to treat widespread penetrability as an empirical claim, though one which is (when properly understood) empirically false.

7. Fodor, surely one of the more austere theorists in this regard, allows for cross-modal effects like the McGurk effect so long as they are domain-specific; the McGurk effect, limited as it is to ambiguous phonemes, is fine in this regard (Fodor 1983: 132fn13).

The need to distinguish merely symbolic anisotropic early visual processes also shapes the constraints on what would count as penetrability. Mere shifts of attention won't count, Pylyshyn thinks, because those can be explained purely in symbolic terms. If I believe the panda is moving, and I redirect my attention, then what I'm *really* doing is something like translating a belief about space back into symbolic terms, and feeding that symbolic representation back into some earlier symbol-using processes. While attention is the most salient escape clause, the more important thing about penetrability—captured by Stokes' clause 2—is that it involves some kind of influence in virtue of the *meaning* of the penetrating relationship. As long as meaning isn't involved, then the division between the intentional and the merely symbolic can be maintained, and the computationalist project survives.

3. The Descending Control of Pain

Return to pain. The architectural discussions have sharpened what's at stake: we need to know whether the effect of cognition on pain is an effect of central states on symbolic processing. The nice thing about pain perception is that the mechanisms for the descending control of pain signals have been mapped out and extensively studied. These pathways provide mechanisms by which the cortex can affect incoming pain signals at the spinal level. I start with a brief review of the peripheral pain system, then focus on the mechanisms for descending control.

The skin and viscera contain a variety of nociceptors that feed information to the pain system. Nociceptors themselves are diverse in nature (Kandel, Schwartz, Jessell, et al. 2000: 473ff.), because possible insults to the body are so various. Some nociceptors are sensitive to ranges of heat or cold. Others are sensitive to mechanical or chemical insult. Many nociceptors are free nerve endings, rather than involving specialized organs. The picture is complicated further by the fact that even peripheral nociceptors are sensitive to context: the so-called *silent* nociceptors in the viscera evoke pain only in the presence of inflammation (McMahon & Koltzenburg 2006: 464). Furthermore, it has long been known that under the right contexts—especially inflammation—ordinarily non-nociceptive receptors like those for touch or stretch can evoke painful responses (McMahon & Koltzenburg 2006: 729ff.).

In addition to differing nociceptors there are also at least two main classes of nociceptive nerve, with myelinated A δ fibers carrying fast information and non-myelinated C fibers slower, longer-term signals. This motley collection of inputs comes together in the dorsal horn of the spinal cord. The laminar organization of the dorsal circuitry is complex enough to perform various kinds of informa-

tion processing before signals are passed on to the brain. A crucial feature of this dorsal circuitry is the so-called *gate control*. First proposed in schematic form by Melzack and Wall (1965), the spinal gate can act as a gain control on inbound signals—a sort of volume knob, if you like. Some inputs (sustained C-fiber firing, for example, or inflammatory compounds) increase the gain. When this goes wrong, one sees phenomena like allodynia, in which otherwise innocuous stimuli like light touch become extremely painful.⁸ Other inputs turn down the volume. Adjacent tactile input fed in by A β fibers often does so, which explains why rubbing an injury can make it hurt less. The view I express is schematic, of course; in a review of the current literature, Todd writes

It is now known that the neuronal organisation and synaptic circuitry of the region are far more complex than could have been imagined at the time of Melzack and Wall's Gate Theory. However, the basic assumption that the superficial dorsal horn modulates nociceptive input is now universally accepted. (Todd 2017: 1)

A further key feature of gate control theory was the assumption of descending spinal pathways to allow for central control of gain. That is, among the inputs that modulate dorsal horn processing are nerves descending from the central nervous system. Central modulation of pain had long been proposed (Fields, Basbaum, & Heinricher 2006: 125). Gate control suggested a concrete mechanism—downward control of gain—by which this could be accomplished.

Several central mechanisms for selective downward modulation have since been elaborated (Fields et al. 2006). Core brainstem structures are the periaqueductal grey (PAG) and the rostral ventromedial medulla (RVM). The PAG is a structure that plays a vital role collecting and integrating information about a variety of stressors, forming an interface between brain and bodily responses (Benarroch 2012). The PAG affects the RVM directly, and the RVM in turn sends descending tracts to the dorsal horn circuitry that is implicated in gate control. The RVM contains both 'ON' and 'OFF' cells that, respectively, enhance and inhibit the dorsal spinal gate. An interesting feature of the RVM circuitry is that it is tonically active—there is always *some* descending control.

Cognitive modulation of pain is affected by a variety of different frontal circuits (Kong et al. 2013). As Fields et al. (2006) note, the connectivity between PAG and the rest of the brain gives straightforward way in which beliefs and desires can influence pain perception directly. The PAG also receives both ascending input from the periphery and downward input from the frontal lobe and amyg-

8. Note that hypersensitivity to touch can probably be driven by a number of distinct phenomena, including entirely receptor-level phenomena as well as spinal and central gating. Thanks to an anonymous reviewer for emphasising this point

dala. Some of these effects are attentional, while others are more direct. For example, they note that “contextual cues acquire the power to either increase or decrease the activity of nociceptive dorsal horn neurons in the absence of a noxious stimulus” through appropriate conditioning, and that this pathway may explain many well-known expectancy effects (Fields et al. 2006: 137).

However, the PAG–RVM axis, while not the only area responsive to or involved in placebo analgesia, appears to be something of a final common pathway through which these influences are funneled (Benedetti et al. 2005; Colloca et al. 2013). There is also direct evidence that the PAG–RVM circuit is a key mechanism in placebo effects. The PAG is one of the classic sites of action of opioids. Levine, Gordon, and Fields (1978) showed that placebo effect is modulated by endogenous opioids using the opioid antagonist naloxone. Functional imaging shows that the effect of naloxone on placebo response is mediated by the PAG–RVM circuit (Eippert et al. 2009a; Linnman et al. 2012). Attentional modulation of pain is also mediated by the PAG (Tracey et al. 2002), as are lesser-known descending inhibitory effects like offset analgesia (Derbyshire & Osborn 2009). Conversely, in rats, activation of on cells in the RVM produces enhanced sensitivity to pain (Neubert, Kincaid, & Heinricher 2004), neuropathic pain is associated with a loss of modulatory cells in the RVM (Leong et al. 2011), and modulating the RVM with lidocaine can both produce or alleviate allodynia depending on context (De Felice et al. 2011). Finally, there is substantial evidence that placebo effects act directly on the spinal circuitry involved in pain. The on and off cells in the RVM project to the dorsal horn of the spinal cord (Fields, Malick, & Burstein 1995). And finally, imaging of placebo analgesia shows that it has direct effects on spinal circuitry (Eippert et al. 2009b).

Return to the question at hand. Given the facts about descending modulation of pain, at what level does that top-down influence appear to be working? It doesn’t appear to be a central effect: placebo analgesia is not a matter of reasoning about what is the case. But nor does it appear to be an influence on some putative modular early processing either. The effects on the brainstem are already much lower down than what we think of as typical cognitive processing. The PAG is a midbrain structure that is otherwise involved in stereotyped behavior such as mating behavior. The medulla is typically associated with basic non-cognitive physiological functions: breathing, heart rate, vomiting, and the like. More importantly, neither seems to be involved in processing *ascending* information about pain: what is affected is not inbound processing.

Instead, the most plausible understanding of downward modulation is (I claim) that it targets the *transducer* for pain. That is, we ought to treat the peripheral systems up to and including the dorsal circuitry as the transducer for pain. Dorsal circuitry is well-posed to serve as part of a transducer. Its function is to make something meaningful from diverse inputs, including both neural and

non-neural information that ranges across different spatial and temporal scales. The most proximal inputs don't mean much on their own: it is only in context that this input means anything relevant to pain. Dorsal circuitry brings together that context and sends on an already-processed signal.⁹ The most peripheral sensory organs convey a lot of noisy information, much of it redundant; sending everything up to the brain would waste bandwidth. Better to process locally and send a consistent signal both up to the brain and back out to local circuits for immediate response.

Further support for the view of the dorsal horn as a computationally complex system comes from quantitative facts about the dorsal horn itself. Nearly all of the pain-related neurons in the superficial dorsal laminae are interneurons, with only a small fraction projecting onward to the brain. Overall there are about half as many inbound afferent inputs as there are interneurons, and about *thirty times* fewer projection neurons.¹⁰ Computationally speaking this setup exhibits a pattern of dimensionality expansion and reduction that is the hallmark of complex computation elsewhere in the nervous system. Furthermore, the interneurons themselves include a mix of excitatory and inhibitory neurons, each of which in turn include of complex and heterogenous subpopulations, creating further opportunity for complex local processing (Browne et al. 2020; Peirs & Seal 2016). Facts such as these lead Braz et al. to conclude that “The dorsal horn of the spinal cord is not merely a ‘relay station’ where primary afferents engage the projection neurons. Rather it is the locus of incredible integration, where sensory information is subjected to local and supraspinally derived excitatory and inhibitory regulation” (Braz et al. 2014: 526).

In conclusion, the mechanisms involved appear to funnel a variety of different influences from across the brain down through a single common pathway. That in turn affects spinal gating mechanisms and hence the inbound

9. In this, we can see a principle that is a fundamental part of most sensory transducers: send only what is informative (Sterling & Laughlin 2015). Many strategies used by sensory transducers are born of the need to make use of tightly constrained bandwidth. By processing information locally and sending on only what is needed, sensory organs allow both for accurate transmission of information and more efficient downstream processing. The dorsal circuitry appears to follow this pattern.

10. The details depend on the species and the location in the spinal cord, but the pattern itself should be broadly applicable. As for specifics: most estimates are done in rats. I have used Chung et al. (1979: 594)'s average of ~6000 afferent axons in the tract of Lissauer for lumbrosacral segments. Schmalbruch (1987) estimates ~12,000 total neurons in the L4 dorsal root. Spike et al. (2003) give a figure of 400 projection neurons from Lamina I of L4. Note that this probably underestimates the magnitude of the expansion stage, as Chung et al.'s estimate covers both myelinated and unmyelinated axons and includes *all* afferents, not just ones relevant to nociception. On the other hand, it appears to be widely accepted that between 5% and 1% of neurons in the superficial laminae project onward (Browne et al. 2020; 2021; Spike et al. 2003; Todd 2010)—so the estimate of the compression stage is, if anything, on the low side.

nociceptive signal. The effect on the inbound signal is at the spinal level, and occurs before the information fans out to the rest of the brain (which it does very widely in the case of pain). That means, however, that the dorsal circuitry would appear to be the primary place where information about insults to the body are translated into a common code and sent on for further analysis by the cognitive system.

Hence we have a possibility that was not anticipated in the original debate (for good reasons, as I'll discuss shortly): downward modulation affects transducers for pain in a systematic, computationally complex way. Is this a sensible position? Is it meaningfully different from thinking about cognitive penetrability as such? I think the answer to both questions is yes; I'll spell out why in the course of responding to some common objections.

4. A Defense of Transducer Calibration

Pylyshyn's original setup left little room for the possibility of effects on transducers, in part because transducers are boring, basic units. As such, one might object to the proposal in several ways, none of which succeed.

4.1. *There Cannot Be Downward Effects on Transducers*

A first objection—rarely offered in so crude a form, but still worth considering—is that our setup does not make any room for downward effects on transducers. Transducers are the starting point (the story goes) of perception, and the flow of information is generally forward, perhaps side-to-side, but mostly one-way. Arguments about cognitive penetrability are about whether late stages of the process can affect slightly earlier ones, and that is contentious. Why think that there could be influence all the way back to the transducer level?

Yet phrased this way, the answer is obvious. *Everyone* thinks there can be influence from central cognition all the way down to the spine and beyond. We call that the motor system. If I desire to drink a sip of coffee and move my hand to grab the cup, my desire has a causal influence that stretches all the way down to the muscles in my arm. The effectors at the end of that process play a structurally analogous role to transducers on the input side of things.

Indeed, though they don't make too much of it, both Fodor and Pylyshn note that their architectural story ought to be symmetrical. That is, once I decide to move my arm, there ought to be some story about how that decision gets translated into purely symbolic terms, that gets processed further, and ultimately bottoms out in effectors at the end. While that end of things hasn't received the same

philosophical attention, there are reasonably worked out, empirically grounded models of how that story might be told (Mylopoulos & Pacherie 2017; 2019).

There is an interesting question about whether this sort of downward effect on the motor system counts as cognitive penetration. My desire for a sip of coffee would seem to affect my arm in something like a ‘semantically coherent’ way, after all, in that moving my arm thus and so is a *way of* satisfying that desire. In an insightful discussion of the motor system, Mylopoulos (2021) suggests that the motor system is open to widespread cognitive penetration for this reason, but there is still substantial evidence for modularity in the sense of informational encapsulation. I suspect that at this point the debate becomes somewhat terminological.

For whatever is going on here, it is commonplace and widespread: it happens any time you have motor action. Presumably that alone is not the sort of penetrability that would threaten the classical architectural picture. So insofar as there is anything striking about the proposal, it is only that one and the same item is implicated as the endpoint of one process and the start of another—the spinal circuitry for pain functions at the same time effector and transducer (in the architectural sense of those terms). There is no obvious reason why things couldn’t be wired up that way; unfamiliarity is no real objection to it. Insofar as there is an objection here, I suspect it is rather to the idea that this influence is complex, and so I turn to that.

4.2. *Transducers Are Primitive, and So Cannot Be Modulated*

Insofar as transducer modulation seems odd, I suspect it is because most people think of transducers as performing a primitive operation. Primitive operations, the objection goes, don’t have internal complexity. So if spinal circuitry is affected, then spinal circuitry is part of the cognitive system proper—because it’s doing something computational, and hence cognitive.

I think this is wrong, and it’s worth unpacking.¹¹ For starters, note that *Pylyshyn* doesn’t just assert that transducers are primitive. Indeed, *Pylyshyn’s* (1984) contains an entire chapter devoted to arguing about transduction. The context is interesting. *Pylyshyn* is concerned with fending off a challenge from Gibson and

11. There is a less interesting terminological form of this objection—most commonly raised to me by electrical engineers and physiologists—that it’s just a matter of convention that ‘transducer’ refers to something primitive and simple. Perhaps, but there are equally good disciplines that allow ‘transducer’ to refer to things with computational structure. So, for example, a *finite state transducer* is a deterministic finite state automata that can translate from strings to strings (Sipser 2013: 87). Insofar as this objection is interesting, it is because it is about more than terminological hygiene.

the ecological psychologists. In the previous sections, I didn't say much about what a transducer was, other than a bridge from the nonsymbolic to the symbolic. The only thing that transducers explain is "how certain nonsymbolic physical events are mapped into certain symbol systems" (1984: 152). That leaves open what sort of "nonsymbolic physical events" can be transduced.

Gibson, infamously, thought that complex relational properties like affordances could be perceived directly (Gibson 2014; see Reed 1996: 64ff. for a succinct introduction). Translated to present terms, Gibson-via-Pylyshyn can be read as saying that you have transducers that are sensitive to (say) grasping affordances, and output a symbol with an *intentional* interpretation 'this is graspable.' That could then enter directly into inferential relationships, skipping the need for the computational, merely symbolic steps entirely.¹² If successful, then this might cut out the need for the intermediate symbolic step altogether—but then you don't get cognitive science, which makes it a nonstarter.

So Pylyshyn asserts three "principal criteria" that a transducer must meet and that are meant to block the Gibsonian move (1984: 153ff.): the function of a transducer is "*primitive* and itself *nonsymbolic*," it is "stimulus-bound" to the environment, and its function must be describable as a mapping between "the language of physics" and "discrete atomic symbols." The paradigmatic transducer is something like a photocell: it says 'o' in the dark, '1' in the light, it doesn't need to do any computing to tell you this, and it doesn't have any choice about the matter.

While austerity serves to make Pylyshyn's point, it is clear that none of these three conditions can be necessary conditions on transducers as such—for there are plenty of things that traditionally count as transducers but that violate all three criteria.

Consider my trusty USB microphone. From the point of view of my computer, it is a transducer: it takes in sound waves and sends on a symbolic representation of that audio stream. Yet the stream is not a series of atomic primitives: it is structured USB-compliant audio.¹³ Nor is it necessarily stimulus-bound—by which Pylyshyn helpfully clarifies as interrupt-driven (1984: 157). There are a number of different USB modes that can be used for data transfer. One mode involves classic interrupts, but others can involve polling by the host computer,

12. This is an idiosyncratic reading of Gibson, given that he has become the champion of the nonsymbolic, antirepresentationalist approach to cognition (Chemero 2009). But this is about Pylyshyn's reading of Gibson; again, if you go antirepresentationalist then you want to do something very different than this paper is doing. Thanks to Tony Chemero for helpful discussion on this point.

13. Pylyshyn later relaxes this to say that the output is either an atomic symbol or an n-tuple (1984: 158ff.) but that strikes me as either still too weak (USB audio streams are not, programmatically speaking, n-tuples) or else it makes the distinction vacuous (because it is true in the sense that *any* structured data can be represented as an n-tuple).

or (as is the case of my microphone) isochronous transfer modes that provides higher bandwidth at the expense of a dedicated handler on the host end.¹⁴

The relationship between the microphone and my computer is also not purely unidirectional. I can control gain, pickup pattern, and mute status using the appropriate software. So in addition to receiving audio from my microphone, I can calibrate it via downward streams of information. Hence there is also a bridge from the symbolic back to the insides. Again, this isn't particularly remarkable: there are plenty of output transducers as well. The fact that a transducer has both an input and an output role should be a conceptually unremarkable combination of the two facts. Finally, my microphone is also computationally complex inside: it is hard work to produce USB audio.

Yet despite that computational complexity, it is—I claim—still obviously a transducer *from the point of view of my computer*. For one, it provides the core function of a transducer: it codes nonsymbolic events in a language the computer can understand. For another, my microphone has a completely different computational architecture than my computer has. It does not have the same CPU. It may not have a CPU at all—it may use special-purpose chips. For all I know, it has vacuum tubes in there. (I checked and it does not, but you get the point.)

Or to put the argument another way: to say that the complexity of the microphone rules it out as a transducer is to say that only the very edge bits in the microphone itself count as transducers, and the rest is part of the computational architecture of my machine. That requires saying extremely counterintuitive things: for example, every time I plug in my microphone the computational architecture of my machine changes in radical ways, because by adding a microphone I add a new set of microphone-specific computational primitives to the system. That is madness. The whole point of standards-compliant accessories is that programmers *don't* have to muck around with their computational ontology in this way: they can leave that to the people who understand the computational demands of microphones, and just work with the nice, well-structured symbols that microphones provide.

Any explanation of a computational system is going to involve a choice about where to draw the line between the computational bits and the non-computational bits. Sometimes this line is obvious. Picking a boundary always involves the same considerations of economy and simplicity and fruitfulness that guide the development of any theory. A computational theory postulates a set of computational primitives—operations and representations—alongside non-computational primitives like transduction. These primitives can be combined to explain the activity of the system. The choice of which computational bits (under some perspective) can be encapsulated and treated as primitives is

14. For details, see (USB 3.0 Promoter Group 2017) particularly §4.4.

really a question about how to build a satisfying computational ontology. The right answer is not always to take the union of everything you find. Sometimes the right answer is to say that there are two distinct computational systems that interact in structured ways via the transfer of physical states that have symbolic status in both (though not necessarily the *same* status).

Note that the USB microphone is, in some sense, the simplest case. It is possible for a transducer itself to have computationally complex transducers as parts: computational systems are often hierarchical. Perhaps at some point we must even bottom out at Pylyshyn's very simple transducers. The point is that 'being a transducer for' is an intransitive relationship: even if there is an eventual bottoming-out, it doesn't mean that the end point is a transducer *from the point of view of my computer*.

Finally, once we admit that transducers can be computationally complex, there does not seem to be a reason to limit transduction to things couched in a mythical 'language of physics.' Indeed, what I pick up might be arbitrarily complex. I want my rail system to avoid hitting pandas. So I buy a Macintosh, train a neural net to recognize pandas, and then use a USB interface to send panda location and bearing to the (architecturally different) computer that controls the train. From the point of view of the train, there is now a primitive transducer that tells it about pandas.

Now, one can admit all of this while still thinking that, when it comes to *cognitive science*, the Gibsonians were over-enthusiastic about the possibilities of transduction. One might notice (for example) that there's a lot of circuitry shared between the transducers for things that Gibson thought you directly perceived, and that this sharing is a good reason to treat that stuff as part of the core computational architecture. There is plenty of conceptual wiggle room between thinking that only the most basic of physical properties are transduced and thinking that *everything* is transduced. But insofar as spinal processing is itself a form of computation, modulated by descending signals, that fact alone does not disqualify it.

4.3. *This Is Just Attention*

A final objection returns to the standard escape clause. Recall that shifts in attention are, by standard convention, accommodated by the traditional story. Biasing inputs in certain ways is common and unremarkable (the story goes): that's just what attention does. So too here: what I've shown is there are some facts about the *mechanism* for biasing inputs, but the traditional story ought to be unimpressed.

I think this is tempting but altogether too hasty. For if 'attention' is to be more than an all-purpose shibboleth in these discussions, the details matter

(Mole 2015). And in this case, I think that there are good reasons to distinguish what's going on from shifts in attention.

For starters, note that shifts in attention are usually *phasic*. That is, there are many ongoing processes, and sometimes we attend to one and sometimes another. We can perfectly well talk about how a process would run in the absence of attention; insofar as attention is a limited resource, we might even consider this the ordinary state of many simple processes. Note that this matches up with standard discussions of cognitive penetration, since it's also usually assumed that penetration is a phasic process. That is, there is an ordinary way that perception might run, and cognition only occasionally intrudes. Bruner and Goodman (1947)'s classic discussion of the effect of value on perception of coin sizes presents an ordinary process that is distorted by need and knowledge. Delk and Fillenbaum (1965) claim that knowledge of an object's characteristic color effects how a neutral cutout is seen has a similar flavor. Again, I emphasize that you don't even need to think that these effects actually *happen*—the point is rather than the architectural logic of cognitive penetrability suggests that it is sensible to talk about ordinary processes of perception onto which cognition occasionally intrudes. And in turn, this is why attention is such a handy escape clause: because insofar as attention is also phasic, it has the right sort of properties to give an alternative explanation.

By contrast, descending modulation of pain transduction is *tonic*. The level of activity of 'on' and 'off' cells in the RVM continually modulates activity in the dorsal horn. That is an important part of the explanation for why many pain syndromes occur (for example): when downward modulation is absent or mistaken, then the transduction of peripheral signals becomes distorted. Conversely, it also explains why painless injury can happen without an initial painful state that is then suppressed: the pain system is *always* being contextually modulated. Hence there is not a straightforward sense in which there is an 'ordinary' course of action for the pain system that is free of such an effect.

Rather than attention, I suggest an alternative functional role for this descending information. There have been numerous specific functional hypotheses about the role of descending control. Yet these functional stories rarely explain why there is a parallel system of *enhancement* of pain, or why the system would appear to have a certain degree of spatial specificity. I think there are two broader explanations to be given, one specific to pain and one more general. Both give a central role to the idea of *calibration*. Roughly speaking, calibration occurs when an input system needs to dynamically adjust its sensitivity based on prevailing conditions.

First, the nociceptive system is (I suspect) unique among sensory systems in one regard: it is designed to deliver accurate information under situations where the *sensory system itself may be damaged*. Furthermore, since injury often comes

along with functional problems, missing or distorted information can itself be relevant to determining the extent of the damage. Conversely, given that many nociceptive inputs also fire in fundamentally innocuous situations—such as deliberately lifting a heavy weight or easing into a hot bath—it can be useful to down-regulate their inputs as well.¹⁵ One plausible role of downward modulatory circuits, therefore, is to selectively calibrate the pain system by dynamically weighting different lines of information.

Our unfortunate builder provides a nice illustration of how this might work. First, nociceptive information from the foot itself would be processed much more slowly than would visual information: the nonmyelinated C-fibers that would carry the bulk of the nociceptive signal from the foot have a conduction velocity of less than 1 m/s (Kandel et al. 2000: 474). Vision is much faster, and so there would be a good second or two during which visual processing would dominate. Given the very plausible inference from vision that there was damage to the foot, this gives a good reason to enhance local processing of ordinarily *non*-nociceptive information that would plausibly be present about the touch and pressure of the nail against the toes. And in other circumstances, that would be a reasonable call: the lack of dedicated nociceptive input might be a consequence of nerve damage. Downward modulation thus provides a mechanism for contextually calibrating the input of the nociceptive system. This can go wrong, of course. It often does. But the need for heavy downward calibration of pain is a consequence of the generally messy circumstances under which pain is called upon to operate.

Second, there is a more general reason why one might want to calibrate pain (or any other input) relatively early on. Pain information enters the brain through several routes, and fans out quickly when it does. That's no surprise: pain is relevant to many, many different cognitive tasks. But pain usually demands a coordinated response. This means that calibration further downstream runs the risk of conflict: if process one thinks that there is a serious problem and process two thinks it is fine, then downstream systems must adjudicate between the two processes or else pick a side (possibly the wrong one) and let it steer the boat. If the inputs to both processes reliably coordinated, on the other hand, the chance of conflict diminishes—and the best way to do *that* is just to calibrate the input from the start.

The converse problem is seen in the other direction. Many processes might have information relevant to the calibration of input. Many of these operate at different timescales. Without some further mechanism for coordination, one can imagine modulation becoming chaotic. Funneling information through the

15. Note that descending control of pain appears to be phylogenetically widespread; see for example Gibbons, Sarlak, and Chittka (2022) for an argument that insects have analogous descending control circuits. This suggests that descending control plays a very important general role.

PAG–RVM gives a mechanism by which these signals can be coordinated. One therefore finds a useful topology for the calibration of pain: a fan-out past the point where modulation occurs, and a fan-in from later stages, through a common pathway, down to the point of modulation.

In sum, we find a functional picture that is very different from one on which attention flexibly and occasionally enhances certain processes. The downward calibration of pain is a mandatory and tonic process.

5. Transducer Calibration

Cognitive penetration is classically presented as the effect of intentional states on sub-intentional processing—mixing memory and desire to stir the dull roots of early vision. The plausible mechanism identified for the downward control of pain, on the other hand, appears to be a mechanism by which intentional states can affect the operation of a transducer, by affecting the activity of something that is effectively treated as a computational primitive from the point of view of central cognition. I have suggested that this downward influence ought to be understood as a distinct process of *transducer calibration*. Like cognitive penetration, it involves (very broadly speaking) the influence of ‘later’ stages of processing on ‘earlier’ ones—indeed, because it is an influence on transduction, the effect of transducer calibration occurs about as early as possible. The topographic features of transducers demand it.

Understood this way, we can also explain why downward influences on pain would appear to have some features that look like penetrability and others that don’t. On the one hand, there is a sense in which transducer calibration needs to be strongly isotropic. Return to painless injury. As I noted above, Wall (1979a; 1979b; 2000) argued that painless injury is adaptive in cases where escape to safety is paramount. Pain tends to limit motion, which works contrary to escape. So it is adaptive to suppress pain until one is safe. The appraisal of personal safety would appear to be a classic case of an isotropic context, because almost any sort of information is potentially relevant: I am probably safe if I am in a hospital—unless I know that the mysterious slasher *is one of the doctors!*

So the *potential* intentional information that can affect transducer calibration is hard to bound. On the other hand, the *structural* facts about the calibration of the pain system—and in particular the fact that the downward effects appear to go through a bottleneck via the PAG and the RVM—suggest that in order to affect calibration, influences must be filtered down through a single, orderly, well-defined point of intervention. Unlike the all-to-all translation problem envisioned by Burnston (2017), then, the orderliness of transducers suggests a straightforward way in which transducer calibration might occur.

This would also explain why you can't just think your way out of pain: while many beliefs and desires might be potentially relevant to transducer calibration, we should expect the downward filtering process to weight these in a survival-appropriate way.

Where does the shift to cognitive calibration leave debates over *cognitive* penetrability? That depends what you care about. If the hope was to use top-down influences on pain to defend classic cognitive penetrability in other—especially visual—domains, then I suspect that this will be something of a letdown: Shevlin and Friesen (2021)'s conclusion that this does not generalize is probably correct.

On the other hand, if you are interested in better picture of the architecture of cognition, then I think the role of transducers remains a rich one ripe for further exploration. I conclude with some speculative reflections on the architecture of cognition and the differentiation between different types of top-down influence.

To start, note that the appeal to cognitive calibration provides explanatory resources that are absent from traditional accounts. So, for example, Casser and Clarke have a brief discussion of influences on the spine, setting them aside “merely eliciting pre-inferential influences” (Casser & Clarke 2022: 14). That is typical of traditional accounts which, following Pylyshyn, ignore transducers. Yet when it comes to the effect of practical assessment on pain, they struggle to find a mechanism by which appraisals of the situation could systematically influence pain. They suggest that there might be other low-level sensory processes that reliably correlate with danger and which could influence pain perception. But they offer no evidence for this, it is difficult to square with other evidence, and does not explain (e.g.) other phenomena like placebo effects or downward inhibition in the absence of danger. So as it stands, the explanation seems more or less *ad hoc*. Yet the need for *ad hoc* explanation stems entirely from the fact that they have tied their hands with respect to the complexity of downward influence on transducers. Once one acknowledges that there are orderly but computationally complex mechanisms by which transducers can modulate pain signals in response to beliefs and desires, one can give more satisfying general explanations that cover a variety of phenomena.

Indeed, while I have focused on pain, I suspect that cases of cognitive calibration might be relatively widespread. Other bodily sensations like hunger and thirst are influenced by anticipation and memory (Kandel et al. 2000: 1007ff.); it is an open question how those effects work. Even the more classic informative modalities might show calibration effects. The eye is adaptively controlled in many ways: where it points, the focus of the lens, and the diameter of the pupil are all at least partially under central control. There is also good evidence that the cochlea has specialized cells that create sound in order to selectively enhance certain frequencies; this is probably under descending control, and the sound produced varies adaptively with task (LeMasurier & Gillespie 2005). Conversely,

failures of calibration often explain otherwise puzzling perceptual phenomena. Tinnitus can be caused when aberrant cochlear sounds overwhelm ordinary ones (Noreña 2015). The IASP now recognizes that many pain syndromes may occur solely due to dysfunctions within the calibration of the pain system itself (Kosek et al. 2016).

Transducer calibration, then, is unlikely to be limited just to pains, or even just to bodily sensations. Furthermore, the more complex the picture of transduction one adopts, the more difficult it becomes to draw the transduction/cognition boundary in the first place. That is not to say that it cannot be drawn, but rather that doing so requires taking into account both computational and topological features of computational processes involved in cognition. We cannot just look to the periphery and call it a day.

Indeed, in a curious passage from Pylyshyn that deserves more attention than it has received, he writes:

As a final remark it might be noted that even a post-perceptual decision process can, with time and repetition, become automatized and cognitively impenetrable, and therefore indistinguishable from the encapsulated visual system. Such automatization creates what I have elsewhere (Pylyshyn 1984) referred to as “compiled transducers.” Compiling complex new transducers is a process by which post-perceptual processing can become part of perception. (Pylyshyn 1999: 360)¹⁶

This is something of a striking admission: it suggests that the automatization of perception via the training of cortical networks *makes entirely new transducers*. Yet this is not implausible either: cortical networks are extremely good at classifying complex inputs—and classification is a natural basis for transduction.¹⁷

Putting aside these complex boundary problems, however, I think that transducer calibration gives us a starting point from which we can usefully taxonomize different types of top-down influence. That is, rather than just divide phenomena into penetration and non-penetration, we can use various topological and computational properties to divide up ways in which cognition might affect other processes.

16. For what it’s worth, I cannot find any reference to compiled transducers in Pylyshyn (1984).

17. For some intriguing remarks in this direction, see Webster’s work on visual adaptation and perception for both color (Webster 2001) and face (Webster & MacLeod 2011) perception. Insofar as these are classically early and late visual processes (respectively), the idea that they might be constantly calibrated by the ambient environment is intriguing. Webster (2001) also argues that Gibson drew the wrong lesson when he turned away from studying perceptual aftereffects; I find a lot to sympathize with in this.

For starters, recall the fan-in structure of information flow involved in calibration. Transducer calibration appears to have a *many-to-one* structure: an indefinitely large set of influences is funneled through a final common pathway in order to calibrate a transducer, which in turn broadcasts widely.

We may distinguish other putative top-down influences in the same way. Shifts in attention are often treated as akin to moving one's head (see, e.g. Siegel 2012: 203). What they have in common is often obscure—but not that both processes involve a *one-to-many* top-down influence: moving your head changes a great number of things, both on the retina and elsewhere. Similarly, if a 'shift in attention' is to be more than an undischarged promise, it must involve a global shift towards some sorts of processing and away from other kinds of processing.

Thinking of attentional effects this way also makes clear that some kinds of competitive processes should *not* be put in the same category, precisely because they seem to have more limited scope. As Mole (2015) points out, there are attentional phenomena that are clearly different in kind. Suppose that attention sometimes does its work via biasing processes that are interacting by competitive inhibition—think of the role of attention in seeing the Necker cube *this* way rather than that, or naming a Stroop color rather than reading the word. These processes *don't* seem akin to merely shifting one's head: instead, they appear to be internal effects in the cognitive system themselves. So the one-to-many criterion seems to do real work in disentangling putative defeaters for penetrability claims.

Finally, one might split claims about classic cognitive penetrability into two kinds. On the one hand, there are *one-to-one* top down effects: beliefs about the color of a banana affect, in a targeted and semantically coherent way, processing of a banana-shape's color. This seems to me to be the sort of influence that Burnston and Cohen (2015) quite rightly point out is conceptually unproblematic if true: we should expect it to be anisotropic if it occurs, and it is mostly an empirical question about whether it does. On the other hand, there might be *many-to-many* top-down effects. Pylyshyn's concern about unrestricted, isotropic top-down influence seems like a concern with this sort of relationship. The picture of top-down influence embodied in recent big-picture predictive coding accounts (Clark 2015; Hohwy 2013) often seems to suggest that there should be this kind of many-many relationship.¹⁸ This is also, of course, the sort of top-down influence that has much more difficulty accounting for *failures* of penetration. And, of course, one-one and many-many are graded notions: there are likely interesting intermediate cases.

18. The degree to which a predictive coding picture must be committed to widespread cognitive penetration, or indeed to any penetration at all, is actually somewhat of a tricky question. Macpherson (2017) provides an excellent review, demonstrating that the presence and degree of penetrability depends on the details of the predictive coding account.

To return to pain and sum up: we began with a question about whether top-down effects on pain ought to count as cognitive penetration. In spelling out the classic picture of penetration, we discovered Pylyshyn's austere picture of transduction. This combined a topological claim (transducers are the maximally peripheral bits of the cognitive system) with an architectural one (transducers are effectively non-computational). The downward calibration of pain showed us that neither claim was necessary, and both were hard to defend. The debates over penetrability have tended to focus on whether an effect is penetrable, looking to see whether a standard set of defeaters for penetrability holds. The present work belongs in a smaller tradition that suggests that these putative defeaters (like transducer modulation) are in fact architecturally interesting and complex enough to be worthy of sustained attention in their own right.

Furthermore, the calibration of pain provides an excellent, empirically grounded illustration of how the many-to-one calibration process might work. It is notable only in the level of empirical attention it has received, however: it is unlikely to be unique. The real lesson of top-down influences on pain, I suggest, is that questions about what is and isn't cognitive penetration cannot be answered piecemeal: they require a broad theory about the architecture of the cognitive system as a whole and the and information flows within it. Pain is not special—but, as ever, it is terribly illustrative.

Acknowledgements

Thanks to audiences at Ruhr-Universität Bochum, The Australian National University, The University of Tübingen, the Italian Association for Cognitive Science, Sydney University, and Australian Catholic University for helpful feedback on previous versions. Thanks to Tony Chemero and Laurenz Casser for particularly helpful discussions. Work on this paper was partly supported by grant TWCF-2020-20539 from the Templeton World Charity Foundation.

References

- Atlas, Lauren Y. (2021). A Social Affective Neuroscience Lens on Placebo Analgesia. *Trends in Cognitive Sciences*, 25(11), 992–1005.
- Beecher, Henry K. (1956). Relationship of Significance of Wound to Pain Experienced. *Journal of the American Medical Association*, 161(17), 1609–13.
- Benarroch, Eduardo E. (2012). Periaqueductal Gray: An Interface for Behavioral Control. *Neurology*, 78(3), 210–17.

- Benedetti, Fabrizio, Helen S. Mayberg, Tor D. Wager, Christian S. Stohler, and Jon-Kar Zubieta (2005). Neurobiological Mechanisms of the Placebo Effect. *Journal of Neuroscience*, 25(45), 10390–402.
- Benedetti, Fabrizio, Michele Lanotte, Leonardo Lopiano, and Luana Colloca (2007). When Words Are Painful: Unraveling the Mechanisms of the Nocebo Effect. *Neuroscience*, 147(2), 260–71.
- Braz, João, Carlos Solorzano, Xidao Wang, and Allan I. Basbaum (2014). Transmitting Pain and Itch Messages: A Contemporary View of the Spinal Cord Circuits that Generate Gate Control. *Neuron*, 82(3), 522–36.
- Browne, Tyler J., Kelly M. Smith, Mark A. Gradwell, Jacqueline A. Iredale, Christopher V. Dayas, Robert J. Callister, . . . , Brett A. Graham (2021). Spinoparabrachial Projection Neurons Form Distinct Classes in the Mouse Dorsal Horn. *Pain*, 162(7), 1977.
- Browne, Tyler J., Mark A. Gradwell, Jacqueline A. Iredale, Jessica F. Maden, Robert J. Callister, David I. Hughes, . . . , Brett A. Graham (2020). Transgenic Cross-Referencing of Inhibitory and Excitatory Interneuron Populations to Dissect Neuronal Heterogeneity in the Dorsal Horn. *Frontiers in Molecular Neuroscience*, 13, 32.
- Bruner, Jerome S. and Cecile C. Goodman (1947). Value and Need As Organizing Factors in Perception. *The Journal of Abnormal and Social Psychology*, 42(1), 33–44.
- Burnston, Daniel C. (2017). Interface Problems in the Explanation of Action. *Philosophical Explorations*, 20(2), 242–58.
- Burnston, Daniel C. and Jonathan Cohen (2015). Perceptual Integration, Modularity, and Cognitive Penetration. In John Zeimbekis and Athanassios Raftopoulos (Eds.), *The Cognitive Penetrability of Perception: New Philosophical Perspectives* (123–40). Oxford University Press.
- Butler, David S. and G. Lorimer Moseley (2013). *Explain Pain* (2nd ed.). Noigroup Publications.
- Casser, Laurenz and Sam Clarke (2022). Is Pain Modular? *Mind And Language*, 38(3), 828–46.
- Chemero, Anthony (2009). *Radical Embodied Cognitive Science*. The MIT Press.
- Chung, Kyungsoon, Lauren A. Langford, Arnold E. Applebaum, and Richard E. Coggeshall (1979). Primary Afferent Fibers in the Tract of Lissauer in the Rat. *Journal of Comparative Neurology*, 184(3), 587–98.
- Clark, Andy (2015). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford University Press.
- Colloca, Luana, Regine Klinger, Herta Flor, and Ulrike Bingel (2013). Placebo Analgesia: Psychological and Neurobiological Mechanisms. *Pain*, 154(4), 511.
- De Felice, Milena, Raul Sanoja, Ruizhong Wang, Louis Vera-Portocarrero, Janice Oyarzo, Tamara King, Michael H. Ossipov, Todd W. Vanderah, Josephine Lai, Gregory O. Dussor, Howard L. Fields, Theodore J. Price, and Frank Porreca. (2011). Engagement of Descending Inhibition from the Rostral Ventromedial Medulla Protects against Chronic Neuropathic Pain. *Pain*, 152(12), 2701–9.
- Delk, John L. and Samuel Fillenbaum (1965). Differences in Perceived Color As a Function of Characteristic Color. *The American Journal of Psychology*, 78(2), 290–93.
- Derbyshire, Stuart W. G. and Jody Osborn (2009). Offset Analgesia is Mediated by Activation in the Region of the Periaqueductal Grey and Rostral Ventromedial Medulla. *Neuroimage*, 47(3), 1002–6.

- Deroy, Ophelia (2013). Object-Sensitivity versus Cognitive Penetrability of Perception. *Philosophical Studies*, 162(1), 87–107.
- Eippert, Falk, Jürgen Finsterbusch, Ulrike Bingel, and Christian Büchel (2009b). Direct Evidence for Spinal Cord Involvement in Placebo Analgesia. *Science*, 326(5951), 404.
- Eippert, Falk, Ulrike Bingel, Eszter D. Schoell, Juliana Yacubian, Regine Klinger, Jürgen Lorenz, and Christian Büchel (2009a). Activation of the Opioidergic Descending Pain Control System Underlies Placebo Analgesia Supplemental Experimental Procedures & Results. *Neuron*, 63, 533–43.
- Fields, Howard L., Allan I. Basbaum, and Mary M. Heinricher (2006). Central Nervous System Mechanisms of Pain Modulation. In Stephen B. McMahon and Martin Koltzenburg (Eds.), *Wall and Melzack's Textbook of Pain* (5th ed., 125–42). Elsevier.
- Fields, Howard L., Amy Malick, and Rami Burstein (1995). Dorsal Horn Projection Targets of ON and OFF Cells in the Rostral Ventromedial Medulla. *Journal of Neurophysiology*, 74(4), 1742–59.
- Firestone, Chaz and Brian J. Scholl (2016). Cognition Does Not Affect Perception: Evaluating the Evidence for “Top-Down” Effects. *Behavioral and Brain Sciences*, 39, 1–77.
- Fisher, J. P., D. T. Hassan, and N. O'Connor (1995). Minerva. *British Medical Journal*, 310, 70.
- Fodor, Jerry A. (1983). *The Modularity of Mind*. MIT Press.
- Gatchel, Robert J., Yuan B. Peng, Madelon L. Peters, Perry N. Fuchs, and Dennis C. Turk (2007). The Biopsychosocial Approach to Chronic Pain: Scientific Advances and Future Directions. *Psychological Bulletin*, 133(4), 581–624.
- Gibbons, Matilda, Sajedeh Sarlak, and Lars Chittka (2022). Descending Control of Nociception in Insects? *Proceedings of the Royal Society B*, 289(1978), 20220599.
- Gibson, James J. (2014). *The Ecological Approach to Visual Perception: Classic Edition*. Psychology Press.
- Gligorov, Nada (2017). Don't Worry, This Will Only Hurt a Bit: The Role of Expectation and Attention in Pain Intensity. *The Monist*, 100(4), 501–13.
- Gligorov, Nada (2018). Telling the Truth About Pain: Informed Consent and the Role of Expectation in Pain Intensity. *AJOB Neuroscience*, 9(3), 173–82. doi:10.1080/21507740.2018.1496163
- Hardcastle, Valerie G. (1997). When a Pain Is Not. *The Journal of Philosophy*, 94(8), 381–409.
- Hohwy, Jakob (2013). *The Predictive Mind*. Oxford University Press.
- Kandel, Eric R., James H. Schwartz, and Thomas M. Jessell. (2000). *Principles of Neural Science* (4th ed.). McGraw-Hill.
- Kong, Jian, Karin Jensen, Rita Loiotile, Alexandra Cheetham, Hsiao-Ying Wey, Ying Tan, . . . , Randy L. Gollub (2013). Functional Connectivity of the Frontoparietal Network Predicts Cognitive Modulation of Pain. *Pain*, 154(3), 459–67.
- Kosek, Eva, Milton Cohen, Ralf Baron, Gerald F. Gebhart, Juan-Antonio Mico, Andrew S. C. Rice, . . . , A. Kathleen Sluka (2016). Do We Need a Third Mechanistic Descriptor for Chronic Pain States? *Pain*, 157(7), 1382–86.
- LeMasurier, Meredith and Peter G. Gillespie (2005). Hair-Cell Mechanotransduction and Cochlear Amplification. *Neuron*, 48(3), 403–15.
- Leong, Mai Lan, Ming Gu, Rebecca Speltz-Paiz, Eleanor I. Stahura, Neli Mottey, Clifford J. Steer, and Martin Wessendorf (2011). Neuronal Loss in the Rostral Ventromedial Medulla in a Rat Model of Neuropathic Pain. *Journal of Neuroscience*, 31(47), 17028–39.

- Levine, Jon D., Newton C. Gordon, and Howard L. Fields (1978). The Mechanism of Placebo Analgesia. *The Lancet*, 312(8091), 654–57.
- Linnman, Clas, Eric A. Moulton, Gabi Barmettler, Lino Becerra, and David Borsook (2012). Neuroimaging of the Periaqueductal Gray: State of the Field. *Neuroimage*, 60(1), 505–22.
- Litt, Mark D. (1988). Self-Efficacy and Perceived Control: Cognitive Mediators of Pain Tolerance. *Journal of Personality and Social Psychology*, 54(1), 149–160.
- Machery, Edouard (2015). Cognitive Penetrability: A No-Progress Report. In John Zeimbekis and Athanassios Raftopoulos (Eds.), *The Cognitive Penetrability of Perception: New Philosophical Perspectives* (59–72). Oxford University Press.
- Macpherson, Fiona (2017). The Relationship between Cognitive Penetration and Predictive Coding. *Consciousness and Cognition*, 47, 6–16.
- McMahon, Stephen B. and Martin Koltzenburg (Eds.) (2006). *Wall and Melzack's Textbook of Pain* (5th ed.). Elsevier.
- Melzack, Ronald and Patrick D. Wall (1965). Pain Mechanisms: A New Theory. *Science*, 150(699), 971–79.
- Melzack, Ronald, Patrick D. Wall, and Tony C. Ty (1982). Acute Pain in an Emergency Clinic: Latency of Onset and Descriptor Patterns Related to Different Injuries. *Pain*, 14(1), 33–43.
- Mole, Christopher (2015). Attention and Cognitive Penetration. In John Zeimbekis and Athanassios Raftopoulos (Eds.), *The Cognitive Penetrability of Perception: New Philosophical Perspectives* (218–35). Oxford University Press.
- Moseley, G. Lorimer (2002). Combined Physiotherapy and Education Is Efficacious for Chronic Low Back Pain. *Australian Journal of Physiotherapy*, 48(4), 297–302.
- Moseley, G. Lorimer and David S. Butler (2015). Fifteen Years of Explaining Pain: The Past, Present, and Future. *The Journal of Pain*, 16(9), 807–13.
- Mylopoulos, Myrto (2021). The Modularity of the Motor System. *Philosophical Explorations*, 24(3), 376–93.
- Mylopoulos, Myrto and Elisabeth Pacherie (2017). Intentions and Motor Representations: The Interface Challenge. *Review of Philosophy and Psychology*, 8(2), 317–36.
- Mylopoulos, Myrto and Elisabeth Pacherie (2019). Intentions: The Dynamic Hierarchical Model Revisited. *Wiley Interdisciplinary Reviews: Cognitive Science*, 10(2), e1481.
- Neubert, Miranda J., Wendy Kincaid, and Mary M. Heinricher (2004). Nociceptive Facilitating Neurons in the Rostral Ventromedial Medulla. *Pain*, 110(1–2), 158–65.
- Noreña, Arnaud J. (2015). Revisiting the Cochlear and Central Mechanisms of Tinnitus and Therapeutic Approaches. *Audiology and Neurotology*, 20(Suppl. 1), 53–59.
- Patterson, David R. (2004). Treating Pain with Hypnosis. *Current Directions in Psychological Science*, 13(6), 252–55.
- Patterson, David R. and Mark P. Jensen (2003). Hypnosis and Clinical Pain. *Psychological Bulletin*, 129(4), 495–521.
- Peirs, Cedric and Rebecca P. Seal (2016). Neural Circuits for Pain: Recent Advances and Current Views. *Science*, 354(6312), 578–84.
- Pylyshyn, Zenon (1999). Is Vision Continuous with Cognition?: The Case for Cognitive Impenetrability of Visual Perception. *Behavioral and Brain Sciences*, 22(3), 341–65.
- Pylyshyn, Zenon W. (1984). *Computation and Cognition*. Cambridge University Press.
- Reed, Edward S. (1996). *Encountering the World: Toward an Ecological Psychology*. Oxford University Press.

- Schmalbruch, Henning (1987). The Number of Neurons in Dorsal Root Ganglia L4–L6 of the Rat. *The Anatomical Record*, 219(3), 315–22.
- Shevlin, Henry and Phoebe Friesen (2021). Pain, Placebo, and Cognitive Penetration. *Mind & Language*, 36, 771–91.
- Siegel, Susanna (2012). Cognitive Penetrability and Perceptual Justification. *Nous*, 46(2), 201–22.
- Sipser, Michael (2013). *Introduction to the Theory of Computation* (3rd ed.). Cengage Learning.
- Spike, Rosemary C., Zita Puskár, David Andrew, and Andrew J. Todd (2003). A Quantitative and Morphological Study of Projection Neurons in Lamina I of the Rat Lumbar Spinal Cord. *European Journal of Neuroscience*, 18(9), 2433–48.
- Staub, Ervin, Bernard Tursky, and Gary E. Schwartz (1971). Self-Control and Predictability: Their Effects on Reactions to Aversive Stimulation. *Journal of Personality and Social Psychology*, 18(2), 157–162.
- Sterling, Peter and Simon Laughlin (2015). *Principles of Neural Design*. MIT Press.
- Stokes, Dustin (2012). Perceiving and Desiring: A New Look at the Cognitive Penetrability of Experience. *Philosophical Studies*, 158(3), 477–92.
- Stokes, Dustin (2013). Cognitive Penetrability of Perception. *Philosophy Compass*, 8(7), 646–63.
- Sullivan, Michael J. L. (2008). Toward a Biopsychomotor Conceptualization of Pain: Implications for Research and Intervention. *The Clinical Journal of Pain*, 24(4), 281–90.
- Thompson, Suzanne C. (1981). Will It Hurt Less If I Can Control It? A Complex Answer to a Simple Question. *Psychological Bulletin*, 90(1), 89–101.
- Todd, Andrew J. (2010). Neuronal Circuitry for Pain Processing in the Dorsal Horn. *Nature Reviews Neuroscience*, 11(12), 823–36.
- Todd, Andrew J. (2017). Identifying Functional Populations among the Interneurons in Laminae I–III of the Spinal Dorsal Horn. *Molecular Pain*, 13, 1–19.
- Tracey, Irene, Alexander Ploghaus, Joseph S. Gati, Stuart Clare, Steve Smith, Ravi S. Menon, and Paul M. Matthews (2002). Imaging Attentional Modulation of Pain in the Periaqueductal Gray in Humans. *Journal of Neuroscience*, 22(7), 2748–52.
- USB 3.0 Promoter Group (2017). Universal Serial Bus 3.2 Specification. Technical Report, USB Implementers Forum, Inc.
- Wall, Patrick D. (1979a). On the Relation of Injury to Pain (The John J. Bonica Lecture). *Pain*, 6(3), 253–64.
- Wall, Patrick D. (1979b). Three Phases of Evil: The Relation of Injury to Pain. *Novartis Foundation Symposium: Brain and Mind*, 69, 293–304.
- Wall, Patrick D. (2000). *Pain: The Science of Suffering*. Columbia University Press.
- Webster, Michael A. (2001). Visual Adaptation and the Relative Nature of Perception. In *Proceedings of the 2001 International Conference on Image Processing* (Vol. 2, 8–11). Symposium conducted at the meeting of the IEEE
- Webster, Michael A. and Donald I. A. MacLeod (2011). Visual Adaptation and Face Perception. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366(1571), 1702–25.
- Wells, Rebecca E. and Ted J. Kaptchuk (2012). To Tell the Truth, the Whole Truth, May Do Patients Harm: The Problem of the Nocebo Effect for Informed Consent. *The American Journal of Bioethics*, 12(3), 22–29.

- Wiech, Katja, Joachim Vandekerckhove, Jonas Zaman, Francis Tuerlinckx, Johan W. S. Vlaeyen, and Irene Tracey (2014). Influence of Prior Information on Pain Involves Biased Perceptual Decision-Making. *Current Biology*, 24(15), R679–81.
- Wiech, Katja, Markus Ploner, and Irene Tracey (2008). Neurocognitive Aspects of Pain Perception. *Trends in Cognitive Sciences*, 12(8), 306–13.