



GRAHEK-STYLE IMPERATIVISM

Abstract: I explore some of the connections between Grahek's model of asymbolic pain, as developed in *Feeling Pain and Being in Pain*, and the contemporary intensionalist discussion over evaluativist and imperativist models of pain. I will sketch a Grahekian version of imperativism that is both true to his main insights and better at confronting some of the challenges that his theory has faced since its publication.

Keywords: pain, imperativism, information processing, asymbolia, Nikola Grahek

1. Introduction

An immensely influential idea in contemporary philosophy of mind and cognitive science is *intentionalism*: the notion that the phenomenal character of experiences is intimately related to (perhaps even the same thing as) their representational content. (Tye 1995; Dretske 2003; Nanay 2010, among many others).¹

Intentionalism is, first, plausible. Focusing for a moment on *perceptual* experiences (seeing a red patch, tasting chocolate, running one's fingers over asphalt), it is very intuitive that the phenomenal character of these experiences is *presentational*: that they (perhaps just) present the world as being one way or another—red here, chocolatey there, bituminous over yonder. One compelling way to capture this presentational character, it would seem, is to make it depend on representational content. The subpersonal-level way of making the same point is to say that, of the very many scientifically significant properties that brain-body-environment dynamics have, their intentional contents are best placed (for example, and among other things, have the best fineness of grain) to describe facts having to do with the phenomenal character of experience.

Second, intentionalism seems like an excellent way to help bridge the explanatory gap: it is one thing to go all the way from (say) neural activity² to

1 In fact, the intentionalist thesis is compatible with the bearers of content not being the experiences themselves (Martínez and Nanay 2024). See also fn. 7.

2 I sometimes use “neural activity” or “brain function”, and sometimes something along the lines of “brain-body-environment dynamics”, to describe the supervenience base for

phenomenology; and another, less daunting, thing to go from representational content to (presentational) phenomenology. The other stretch of the gap, from neural activity to representational content, also looks tractable, with a few promising research programs in naturalistic metasemantics happily forging ahead (Millikan 1984; Dretske 1988; Skyrms 2010; Neander 2017; Shea 2018; Martínez 2019).

On the other hand, perceptual experience is not the end-all and be-all of phenomenology. Pain (and thirst, thermal (dis)pleasure, orgasms, affective phenomenology in general) are just as important as perceptions to our mental life, if not more. Yet it feels as much more of a stretch to think of pain-related affective phenomenology as somehow presentational: the characteristic way in which pains hurt do not seem to present the world in some way or another—pain just hurts.

One alternative that has been explored at this juncture is, first, to substitute “presentational” by something along the lines of “motivational”, or “pushy”: the characteristic way in which pains hurt may not strike us as presentational, but it (hopefully) does strike us as motivational and pushy. And, second, by developing an imperativist alternative to representational content: a brand of (motivational, pushy) intentional contents that instead of conveying something like correctness conditions (such as, e.g., “there’s chocolatey stuff over here”) convey something like satisfaction conditions (e.g. “fix that issue in your leg!”). Luckily, the research programs in metasemantics I alluded to above, e.g., Millikan’s (1984) ‘biosemantics’, or Shea’s (2018) ‘varitel’ semantics, recognize states with *imperative* contents alongside the more traditional representational, or *indicative*, ones. Millikan calls them “imperative intentional icons”; Shea, “directive representations”.³

Nikola Grahek was not directly interested in what could now be described as the “indicativist-imperativist debate” among intentionalists about pain: the question of whether pain (only) has a presentational, perhaps evaluative content along the lines of “What is currently happening to your right arm is bad news” (Tye 2006; Bain 2013) or whether it (also) has an imperative content along the lines of “Do something about your right arm!” (Klein 2007, 2015a; Martínez 2011; Barlassina and Hayward 2019). He couldn’t have been: he died in 2003, and the first explicit *imperativist* treatment of pain (i.e., one along the lines sketched above) was Klein (2007). Still, much of the discussion in his remarkable *Feeling Pain and Being in Pain* (Grahek 2007, *FPBP* henceforth) is relevant to, and indeed powerfully influenced, this debate. And, I think, the opposite direction is equally interesting: the way of

experiences so as not to prejudice any questions about classical or, e.g., radical embodied (Chemero 2011) cognitive science as the right ways to approach these issues. For what it’s worth, I believe representationalist and 4E approaches to cogsci are all useful, largely compatible and complementary, but nothing of what I will be saying here depends on this.

- 3 While I prefer to think in terms of imperative and indicative *contents*, some readers will rather endorse a model of forceless contents with indicative or imperative forces attached to them. Most of what I will have to say is compatible with both alternatives.

looking at pain that Grahek pioneered could probably benefit from some of the nuance that the above debate has uncovered and developed.

Here I will explore some of these possible connections. First, in §2, I take up Grahek's distinction between "subjectivist" and "objectivist" approaches to painful phenomenology. I will suggest that, the way intentionalist positions are usually developed, one can be both at the same time. Then, in §3, I discuss Grahek's incomplete-information model of asymbolic pain processing. In §4 I offer my own attempt at bringing this model up to date, and patching up some of its deficiencies. In §5 I offer some concluding remarks.

2. Grahek-Objectivism and Grahek-Subjectivism

In FPBP, Grahek distinguishes between *subjectivist* and *objectivist* stances about pain. For the subjectivist, "the sensation of pain with its distinctive phenomenal content or quality ... is the essential component of our total pain experience and plays the central or fundamental role in it" (p. 76, all page or chapter numbers without reference correspond to FPBP). For the objectivist "the feeling of pain is just the awareness of objective bodily states of affairs: the perception or sensory representation of bodily or tissue damage." (p. 78).

Grahek appears to think that these two positions are at odds with one another; but a contemporary intentionalist could perfectly well claim (with Grahek's subjectivist) that the essential component of pain is its phenomenal character, *and* also claim (with Grahek's objectivist) that this phenomenal character boils somehow down to a representation of bodily damage. A coherent way to do this is to endorse the identification or supervenience of presentational phenomenology on representational content, as I briefly discussed above. Another option, which Grahek again couldn't have considered, is to conjoin his subjectivism with an "imperativist objectivism" that makes phenomenal character depend on an imperative intentional state.

It is telling that Grahek sees objectivism and subjectivism as dichotomous. This plausibly stems from his belief that the qualitative component of pain is constitutively independent of any and all epistemic or behavioral consequences that pain may have: "the pure juice or essence of pain experience ... [turns] out to be a blunt, fleshless, inert sensation pointing to nothing beyond itself, leaving no traces in the memory and powerless to move the body and mind in any way" (p. 76). Grahek presents this claim as supported by asymbolia experimental results. In a nutshell: given that asymbolics report that they feel pain, yet do not seem to care about bodily integrity, it must be because pain is inert in the way he so eloquently describes in the quote above. But, of course, there are more than a few interpretations of these results. One possible, and indeed natural, one is that asymbolics are undergoing a state the phenomenal character of which they identify as pain, but that, precisely because it is "powerless to move the body and mind", lacks an important phenomenal

component—the *lived significance* of typical, painful pain. That pain has a lived significance, an affective, action-informing phenomenology is indeed one of the starting points of imperativism. For “hybrid imperativists” such as Martínez (2011), it is one ingredient in the overall phenomenological profile of pain, alongside a more sensory-like, “there’s something going on in my leg” indicative phenomenology that survives in asymbolia. For others, such as Klein (2007, 2015a) this kind of imperative lived significance exhausts the phenomenology of pain. Grahek’s decision to take asymbolic reports at face value was epistemically risky: asymbolic patients, and indeed everyone else, have only had nociceptive sensory (presentational) experiences in the context of pain episodes. It makes perfect sense that they would identify them as pain. But this doesn’t mean that that’s all pain is, and doesn’t mean that therefore pain is blunt and fleshless.⁴

A Grahek-subjectivist could coherently insist that they take phenomenology to be an (or the) essential component of pain, while adding that, over and above the phenomenological ingredient that survives in asymbolic mental life, there is another one: lived, or qualitative, painfulness. Now, to be frank, I am not entirely sure how Grahek thought of phenomenological painfulness. He does make an allusion to “how one feels (affectively, not sensorially)” (p. 77), which looks like a recognition of the existence of affective phenomenology, but this comes right after a discussion of how “the quality of the sensations that one feels is quite irrelevant for determining whether one is in pain”, which implies that *all* of the feelings associated with pain are irrelevant to the role pain plays in our mental economy. Grahek (1991) saw affectivity as a purely attitudinal state, that might be theorized as a replacement of “some irreducible felt quality” (Grahek 1991: 249). At this point in his thinking, it would seem, affectivity was not qualitative—but this is almost a decade before FPBP. So, I don’t know. The idea of affective phenomenology has been recognized in philosophical discourse for a long time, but it is not impossible that it has gained wider currency in recent years, and anyway since Grahek’s book was written.

Grahek is famous for his forceful defense of pain as a complex phenomenon, with at least “sensory-discriminative”, “emotional-cognitive”, and “behavioral” components (p. 2). What I am suggesting here is that it is possible, *pace* Grahek, to think of this complexity as being fundamentally phenomenal. More carefully: there is no obligation on the part of the objectivist to claim that painful phenomenology is the sole essence of pain, although this option is still open to them. There are good reasons to think that pain is indeed a complex phenomenon, not exhausted by its phenomenal profile (Casser 2020; Corns 2020; Serrahima and Martínez 2023). But if we think, apparently with Grahek, that “complexity” requires that pain have both phenomenal and non-phenomenal components, then it is less clear that asymbolic reports are among these reasons.

4 Two recent elaborations of the idea that asymbolic pain is not actually pain are Park (2023) and Griffith and Kind (2023).

Leaving now aside whether we should think of the affective component of pain as phenomenal or not, I turn to the question of how we should think of the way the presentational-indicative, and the pushy-imperative ingredients of pain are organized in our mental life.

3. The Contents of Pain

Pain is (among other things, but centrally) a personal-level phenomenon: *we* hurt, and not just some subpersonal system or other. Even if Grahek paid more attention than most to the neural underpinnings of his target phenomenon, he consistently prefers personal-level analyses in his book. For example, he argues that painful phenomenology cannot represent *tissue damage* (the way evaluativists such as Tye or Bain, cited above, would want) because, if it did, asymbolics would be motivated to do something about it, but instead they just shrug it off:

[T]he pain that these patients feel does not represent for them any damage or potential damage to their bodies. That this is so is best proved by the fact that they consistently smile or laugh during pain testing procedures. (p. 80)

Klein (2015b: 508) responds that representation of damage is not necessarily the right kind of entity to generate motivation. Representations (that is to say, *indicative* intentional states) are, primarily, in the business of *informing* their recipient (be it a person, or a subpersonal cognitive mechanism) about their content. They are in the business of *presenting* their content, as I put it above. Sticking for the moment to personal-level, psychological explanations, an inferential step seems to be needed to go from information to motivation (“I have sustained damage in my foot. That’s not good. Better do something about it!”), and there’s no guarantee that this step will be taken. Klein (2015b) offers one reason why it might not: the owner of the damaged body might not care about it very much (“I have sustained damage in my foot. So what.”). Klein takes this to be an explanation of the shocking behavior of asymbolics when in pain: they do feel it, but seeing as the intentional object of pain (that is to say, the damaged body) is irrelevant to them, they gleefully disregard it. Suppose that you received an SMS informing you that a bicycle you know not or care not about has been stolen or damaged. You’d be likely to gleefully disregard it as well. For asymbolics it is *their* bicycle and they should care—they just don’t.

Klein’s explanation of asymbolic behavior (like Grahek’s) relies on personal-level, rationality-based considerations: given that asymbolics don’t care about their own body, it is rational for them to simply shrug off, or even laugh at, advice and requests having to do with its welfare. There is, though, one shortcoming of this approach to asymbolia: it overlooks how disrupting *the very presence of pain signals*, of the vehicles themselves, can be, regardless

of whether they are semantically pertinent or not—whether, that is, we care about their content. It is one thing to receive the odd SMS about random stolen bicycles, and another thing to be flooded with stolen-bicycle spam (Martínez 2015, 2022): vehicle properties (no pun intended) are relevant, and vehicles, such as pain signals, that can consistently promote themselves to the top of the agent's to-do list (Klein and Martínez 2018) can be very disrupting even if we don't care about their content. For another example, think of earworms: they can make it hard to focus on other things, and can generally be mildly disrupting, regardless of the semantic content (if it has it) of the song that you cannot currently shake off. Asymbolic pain should be at the very least as disrupting as “Total Eclipse of the Heart” going insistently through your mind while you try to engage your attention in other tasks; yet Klein's depersonalization-based model does not seem to speak to how the asymbolic is able to effortlessly ignore the spammy repetition of pain vehicles. An explanation of this fact does not obviously fall off the fact that asymbolics are not interested in the integrity of their body. This is not a blanket rejection of Klein's interpretation of asymbolia, by the way. Asymbolic behavior is very unusual and surely depends on very unusual cognitive substrates. Abnormal disregard for the welfare of one's own body can certainly be relevant to explaining asymbolia—without being the whole story; and rationality-based considerations can be an excellent guide in our theorizing about asymbolia, also without being the whole story.

So here at least personal-level analyses do not seem to be fully adequate. Grahek also relies on an algorithmic-level model of asymbolia, a version of which is better placed to guide inquiry as to the contents of pain-related representations. Grahek endorses the view of asymbolic pain as the result of a disconnection syndrome (Berthier, Starkstein, and Leiguarda 1988): a lesion not in “primary motor, sensory and limbic areas” but in the “association cortices” connecting them (Catani and Ffytche 2005: 2224). In particular, asymbolia would be a sensory-limbic-disconnection problem (chapter 5): Grahek's main intuition here is that if “somatosensory cortical areas” and “limbic structures” are not adequately connected then the asymbolic subject can “recognize the modality, qualities, intensity, and location of noxious stimuli” but (because links to the limbic system are severed) fail to “attach appropriate emotional significance” to such stimuli (p. 52).

The way Grahek presents this “incomplete information” model, though, there appear to be few constraints on information combination: in asymbolic pain, limbic structures (because they are disconnected from sensory cortices) are unable to suffuse the perception of pain with emotional significance; this suggests that such structures are the proprietors of emotional significance, and have (in normal cases) the power to sprinkle (whatever?) sensory experiences with it. This is unsatisfactory, if only because the proposed pain-processing architecture looks suspiciously coarse-grained: some brain area takes care of deciding whether some sensory event counts as pain; some other brain area

takes care of deciding whether that sensory event should present a negative affective profile; if these two areas are adequately connected, normal pain happens; if they fail to make contact with one another, asymbolia happens.⁵

There is an implicit assumption in the Grahekian proposal for a pain-processing architecture that those sensory and limbic processing “boxes” correspond to personal-level features of pain: the “sensory” aspect of pain is the pain quale, the above “pure juice ... of pain experience” as the person in pain feels it; while the “limbic” aspect tracks whether the person in pain cares about it or not. The contemporary shift towards complex-system thinking in neuroscience (whereby whole populations of neurons interact to modulate aspects of brain dynamics that eventually result in behavior) makes Grahek’s two-node (sensory + limbic) network appear somewhat simplistic. Again, this doesn’t mean that simple, neat, human-readable mappings between neural activity and personal-level psychology never happen. They do. But, I submit, it does mean that substantial evidence should be produced for the mapping proposed. I don’t think that Grahek’s modeling exercise, suggestive as it undoubtedly is, provides such evidence.

4. A Better Model of Information Combination, and Back to Imperativism

Still, I think that Grahek’s incomplete-information model is on the right track. If we zoom out enough, cognitive systems do encode sensory stimuli in a way that can subsequently be decoded as (a contribution to) motor representations or behavior (Kriegeskorte and Douglas 2019). This is not a very substantial claim, either: the claim that brain activity encodes information about the world as perceived, and that it is decodable as contributions to motor representations, just *is* the claim that behavior is not probabilistically independent of what the current circumstances of the agent are like. Agents do not behave at random, but in a way that is sensitive to their whereabouts, and their current goals and plans.

In this context, it should be pointed out that the obvious algorithmic choice is to compute sensory-pain status first, and to process “emotional significance”, or affectivity next: presumably, much, if not all, of what goes into the former computation (say, the amount of mechanical stimulation that a certain body is receiving, or its surface temperature, and how much is enough for potential damage) is relevant to the latter computation too. If, e.g. your arm is being compressed, or heated, to an extent that risks tissue damage, this should result in sensory pain *and* should also result in that pain being affectively charged to some extent—so that you feel compelled to do

5 Klein (2015b: 500) also points out that disconnection syndromes seem to depend on “an entirely serial, feed-forward picture of the brain”. But one could read (or perhaps slightly amend) Grahek as discussing abstract information-processing stages, not necessarily related in a neat, one-to-one fashion to brain areas. That’s anyway how I intend the discussion in §4 to be read.

something about it. There are contexts in which a certain typically-painful event should pass relatively unnoticed; for example if the agent has other more pressing concerns. But, even in those cases, sensory processing is centrally relevant to the calculation of whether the agent should be presented with a typical, negatively valenced pain sensation. This is very clear in homeostatic sensations, such as thirst or hunger, where information about current levels of satiety and the current availability of food and drink are combined in the generation of more or less peremptory sensations (Juechems and Summerfield 2019). The same happens with thermal discomfort and pain (Mower 1976).

A more apt way of thinking of the relation between sensory pain and emotional significance than the Grahekian picture of an emotional layer on top of inert sensory juice is, then, to see the latter as an informational input to the former. The picture would be something like this (Martínez and Klein 2016; Martínez and Barlassina forthcoming). We can think of cognitive processing as the progressive coming into focus of what to do next, starting from how things are. Roughly, the closer a state is to the input side of cognitive processing, the more information it carries about objective, sensory-like aspects of the world: the more indicative-presentational it is. The closer it gets to the output side, the more information it carries about what to do next: the more imperative, “significant”, or motivational it is. Sensory pain and lived (or attitudinal) significance would be close to consecutive stages in this causal process of generating behavior from sensation. Pain has indicative and imperative components not because different modules compute them and then some third entity connects them together, but because *there is a gap that needs to be bridged, from having information about potential bodily damage to issuing a command to prevent or fix it*. Almost universally, information about potential bodily damage will have a significant impact on behavior and intentions—a plan to deal with that damage (if it has already happened) or to prevent it (if not) will almost universally be promoted to higher positions in the agent’s to-do list. To that extent, the merely informative component of pain will also have imperative (and therefore affective) overtones. These imperative-affective components will be more central to the subject’s pain experience the more catastrophic the (represented) damage is.

This provides a very natural way of reading the research by Price (2000), cited in FPBP (chap. 5). According to Price, first, the causal direction of pain processing is always from sensory to affective (Price 2000: 1769): affective appraisal is a result of previous sensory processing, and never vice versa. And, second, the affective dimension of pain seems to result from the integration of “somatosensory nociceptive input with other contextual inputs to provide an overall sense of intrusion and threat to the physical body and self” (Price 2000: 1771). That affective phenomenology goes hand in hand with information integration can be seen as a more concrete elaboration of my “coming into focus” metaphor above. Even if mere nociceptive information is relevant to behavior, a more complete assessment of how much of a threat a certain noxious stimulus represents is more relevant. This integrative stage,

the idea would go, marks the crossover from mainly indicative-presentational to mainly imperative-affective.

Grahek's claim that it is limbic processing that "attaches emotional significance" to pain conjures a potentially misleading picture of limbic structures simply making the subject in pain care about their pain—by some sort of stipulative *fiat*, as it were. A more plausible (but still broadly Grahekian) understanding of the involvement of the limbic system in affectivity is as the locus of informational combination of nociception with other goal-relevant and behavior-relevant information, in the process of deciding whether issuing an imperative for the subject in pain to fix the underlying damage is the pertinent thing to do.

The evidence pointing to lesions on the way from somatosensory to limbic cortices, or the limbic cortex itself, as the main kind of brain insult underlying asymbolia (BSPS, chap. 5, Price 2000) makes perfect sense under this interpretation: if it is the limbic part of the pain matrix that is in charge of making the final assessment as to the urgency and importance of a bodily threat, and those areas no longer exist, it's not surprising that we witness the main symptom of asymbolia: a perception-like (i.e., informative, indicative) representation of bodily damage, Grahek's "pure juice or essence", without the "fix that damage!" imperative that limbic structures compute, and which would underlie Grahek's "emotional significance" (or alternatively a painfulness, lived-significance *quale*, depending on whether we take the affective component of pain to be attitudinal or phenomenal.)

5. Conclusion

Grahek is committed to the claims that pain is a complex phenomenon, *and* that asymbolia supports this view. Here I have sketched a package of views that vindicates these two claims—or close enough:

- The complexity of pain is, at least partly, phenomenal complexity:⁶ pain has sensory-like phenomenology; and it also has affective phenomenology. As Grahek's subjectivist would want.
- As Grahek's objectivist would want, these two kinds of phenomenology depend on intentional content, of two different kinds: indicative contents for sensory phenomenology, imperative contents for affective phenomenology.⁷
- Indicativeness and imperativeness, in turn, at least partly depend on the processing stage at which the relevant state happens. When states

6 At least the kind of complexity that Grahek cares about, which doesn't necessarily exhaust the kinds of complexity that, e.g., Corns (2020) or Hardcastle (1999) care about.

7 The way I have described things, pain has two phenomenal aspects, each of which corresponds one-to-one to a kind of intentional state. A, perhaps more apt, alternative is to think of pain experiences as unified, yet dependent on various different representations scattered through a certain cognitive system (Martínez and Nanay 2024).

are closer to the sensory periphery, they are typically more informative about the world as it is perceived than about what the perceiving subject will do next. These states are predominantly indicative. As sensory information gets combined with other sources of information (about goals, the inner state of the subject and their context, etc.) sensory details get progressively unfocused and action-relevant information progressively focused—that is what imperativalness consists in.

- The pain-processing sketch according to which pain starts (sometimes) as an episode of nociception and then gets integrated, in the limbic system, with other subject-relevant information into a more accurate assessment of threat to body integrity is an excellent fit with the above indicativeness-to-imperativalness progression: roughly, nociception would be indicative and therefore sensory-like; the result of limbic integration, imperatival and therefore affective.
- If the structures in charge of issuing the painful imperative are not there, or are cut off from sensory cortices, we should expect the familiar asymbolic pattern of sensory pain without painfulness.

All of this is, of course, programmatic, and each ingredient of the above sketch would benefit from elaboration—a first stab at which has already been taken in some of the papers cited above. It also leaves several important open questions: how should we think of asymbolic *chronic* pain, for example? But, I submit, it paints a preliminary picture that is both empirically well motivated, and can be seen as fleshing out some of Nikola Grahek's main insights. We could call it *Grahek-style imperativism*.

Acknowledgments

I would like to thank Carlota Serrahima for comments on an earlier draft, and Voin Milevski for his helpfulness and patience.

Funding

Financial support was provided by the Spanish Ministry of Science and Innovation, through grants [PID2021-127046NA-I00] and [CEX2021-001169-M (MCIN/AEI/10.13039/501100011033)]; and by the Generalitat de Catalunya, through grant [2021-SGR-00276].

References

- Bain, David (2013) 'What Makes Pains Unpleasant?', *Philosophical Studies* **166**: 69–89. doi:10.1007/s11098-012-0049-7.
- Barlassina, Luca and Max Khan Hayward (2019) 'More of Me! Less of Me!: Reflexive Imperativism about Affective Phenomenal Character', *Mind*. doi:10.1093/mind/fzz035.

- Berthier, Marchcelo, Sergio Starkstein, and Ramon Leiguarda (1988) 'Asymbolia for Pain: A Sensory-Limbic Disconnection Syndrome', *Annals of Neurology* **24**: 41–49. doi:10.1002/ana.410240109.
- Casser, Laurenz C (2020) 'The Function of Pain', *Australasian Journal of Philosophy* **0**: 1–15. doi:10.1080/00048402.2020.1735459.
- Catani, Marco and Dominic H Ffytche (2005) 'The Rises and Falls of Disconnection Syndromes', *Brain* **128**: 2224–39. doi:10.1093/brain/awh622.
- Chemero, Anthony (2011) *Radical Embodied Cognitive Science*. MIT press.
- Corns, Jennifer (2020) *The Complex Reality of Pain*. Routledge.
- Dretske, Fred (1988) *Explaining Behavior. Reasons in a World of Causes*. The MIT Press.
- Dretske, Fred (2003) 'Experience as Representation', *Philosophical Issues* **13**: 67–82. doi:10.1111/1533-6077.00005.
- Grahek, Nikola (1991) 'Objective and Subjective Aspects of Pain', *Philosophical Psychology* **4**: 249–66. doi:10.1080/09515089108573029.
- Grahek, Nikola (2007) *Feeling Pain and Being in Pain*. MIT Press.
- Griffith, Trevor and Adrian Kind (2023) 'Pain Asymbolia Is Not Pain', *Philosophy of Science* 1–18. doi:10.1017/psa.2023.167.
- Hardcastle, V G (1999) *The Myth of Pain*. The MIT Press.
- Juechems, Keno and Christopher Summerfield (2019) 'Where Does Value Come From?', *Trends in Cognitive Sciences* **23**: 836–50. doi:10.1016/j.tics.2019.07.012.
- Klein, Colin (2007) 'An Imperative Theory of Pain', *The Journal of Philosophy* **104**: 517–32.
- Klein, Colin (2015a) *What the Body Commands: The Imperative Theory of Pain*. MIT Press.
- Klein, Colin (2015b) 'What Pain Asymbolia Really Shows', *Mind* **124**: 493–516. doi:10.1093/mind/fzu185.
- Klein, Colin and Manolo Martínez (2018) 'Imperativism and Pain Intensity', in David Bain, Michael Brady and Jennifer Corns, eds., *Philosophy of Pain: Unpleasantness, Emotion, and Deviance*: 13–26. Routledge. <https://www.routledge.com/Philosophy-of-Pain-Unpleasantness-Emotion-and-Deviance-1st-Edition/Bain-Brady-Corns/p/book/9780815361640>.
- Kriegeskorte, Nikolaus and Pamela K Douglas (2019) 'Interpreting Encoding and Decoding Models', *Current Opinion in Neurobiology* **55**: 167–79. doi:10.1016/j.conb.2019.04.002.
- Martínez, Manolo (2011) 'Imperative Content and the Painfulness of Pain', *Phenomenology and the Cognitive Sciences* **10**: 67–90.
- Martínez, Manolo (2015) 'Pains as Reasons', *Philosophical Studies* **172**: 2261–74. doi:10.1007/s11098-014-0408-7.

- Martínez, Manolo (2019) 'Representations Are Rate-Distortion Sweet Spots', *Philosophy of Science* **86**: 1214–26. doi:10.1086/705493.
- Martínez, Manolo (2022) 'Imperative Transparency', *Mind* **131**: 585–601.
- Martínez, Manolo and Luca Barlassina (forthcoming) 'The Informational Profile of Valence: The Metasemantic Argument for Imperativism', *British Journal for the Philosophy of Science*. doi:10.1086/726998.
- Martínez, Manolo and Colin Klein (2016) 'Pain Signals Are Predominantly Imperative', *Biology & Philosophy* **31**: 283–98. doi:10.1007/s10539–015–9514-y.
- Martínez, Manolo and Bence Nanay (2024) 'Many-to-One Intentionalism', *The Journal of Philosophy* **121**: 89–107. doi:10.5840/jphil202412125.
- Millikan, Ruth Garrett (1984) *Language, Thought and Other Biological Categories*. The MIT Press.
- Mower, George D (1976) 'Perceived Intensity of Peripheral Thermal Stimuli Is Independent of Internal Body Temperature', *Journal of Comparative and Physiological Psychology* **90**: 1152–55. doi:10.1037/h0077284.
- Nanay, Bence (2010) 'Attention and Perceptual Content', *Analysis* **70**: 263–70. doi:10.1093/analys/anp165.
- Neander, Karen (2017) *A Mark of the Mental: In Defense of Informational Teleosemantics*. MIT Press.
- Park, Thomas (2023) 'On the Alleged Evidence for Non-Unpleasant Pains', *Inquiry* **66**: 738–56. doi:10.1080/0020174X.2019.1658625.
- Price, Donald D (2000) 'Psychological and Neural Mechanisms of the Affective Dimension of Pain', *Science* **288**: 1769–72. doi:10.1126/science.288.5472.1769.
- Serrahima, Carlota and Manolo Martínez (2023) 'The Experience of Dysmenorrhea', *Synthese* **201**: 173. doi:10.1007/s11229–023–04148–9.
- Shea, Nicholas (2018) *Representation in Cognitive Science*. Oxford University Press.
- Skyrms, Brian (2010) *Signals: Evolution, Learning & Information*. New York: Oxford University Press.
- Tye, M (1995) *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. The MIT Press. Bradford Books.
- Tye, M (2006) 'Another Look At Representationalism About Pain', in M Aydede, ed., *Pain: New Essays on Its Nature and the Methodology of Its Study*: 99–120. The MIT Press. Bradford Books.