

October 2023

Forthcoming in the *Belgrade Philosophical Annual* (special issue honoring Nikola Grahek)

Is Pain Representational?

Murat Aydede

University of British Columbia

Department of Philosophy

`murat.aydede@ubc.ca`

Representationalism in philosophy of perception has become more or less the dominant view. There are various versions of it not all of which are motivated by the same set of concerns. Different metaphysical and epistemological agendas are at work in different strands of the movement. In this paper, I will focus on what has come to be known as strong representationalism. This view has reductive and non-reductive versions, which are usually paired with realist and irrealist versions respectively. Here, first, I will develop a simple, largely empirical, argument against realist reductive version. Later, rather more briefly, I will extend the argument to cover irrealist representationalism.

The main tenet of strong representationalism is that the phenomenal character of any conscious (sensory-affective) experience is metaphysically determined solely by what is represented by that experience, its representational content. This view implies that there is no sensory-affective experience that is not representational in this way. More precisely, it implies that there is no phenomenal aspect of an experience that cannot be reduced back to something that the experience represents. Strong representationalism thus becomes a metaphysical program when it's combined with a research program that seeks to explain the facts that determine how any experience can have representational content in completely naturalistic terms, i.e., when combined with a naturalistic account of experiential representation. The goal and the promise of the reductive strong representationalism is the complete naturalistic (physicalistically kosher) explanation of experiential consciousness. Reductive representationalists also tend to be realist about the sensible properties that are (non-conceptually) represented by sensory-affective experiences. These sensible properties may be complex and gerrymandered (or, even relational) but they are physical properties in good-standing whose instances exist independently of any sensing (representing) mind.

Strong representationalism also implies a strong transparency thesis about experiences: any aspect of an experience that is phenomenally accessible to the experiencer is an aspect of

the content of that experience, an aspect of what is represented — not an aspect of the experience *qua* representor.¹

Realist Representationalism

I believe that reductive realist strong representationalism is false on the ground that there are sensory-affective experiences, such as sensory pains, itches, tickles, orgasms, most sensory pleasures and many more unnamed sensations and bodily feelings, that are not strongly representational in the way required by the view.

Here is the **Ur-argument** for this claim:

(1) For every sensory modality (or submodality) *m* that gives rise to conscious experience *e* in that modality, there is a set of low-level sensible properties *SP* associated with *m* such that *e* represents *SP* only if *e* is the immediate epistemic basis for the acquisition and deployment of the *proprietary concepts* of *SP*.

(Definition)

c is a *proprietary concept* of *SP* IFF

- (i) *c* is needed for the explanation of how the agent *g* (where *c* belongs to *g*'s set of mental representations) discriminates or learns to discriminate instances of *SP*, and
- (ii) *c* is necessarily acquired from sensory experiences *e* in *m*, and
- (iii) *c* has **labeling** uses in that they are non-inferentially (directly/immediately) applied by *g* to the instances of *SP* on the basis of *e* — thus generating true or false *de re* thoughts non-inferentially based on veridical or non-veridical sensory episodes.

(2) None of the candidate proprietary concepts for sensory pain, itch, tickle, orgasm experiences have labeling uses.

(3) Thus, there are no proprietary concepts for these experiences.

(4) Thus, these experiences do not represent any *SP* associated with nociceptive modality and other relevant sensory modalities or submodalities.

(5) Reductive realist strong representationalism is false.

¹ See my (2019a) for a similar criticism of reductive representationalist transparency claims. For the purposes of this paper, I am leaving *imperativism* about pain aside. Its defenders claim that pain experiences are strongly intentional if not representational. The idea is that pain experiences don't have truth-apt representational content: rather, they have strongly intentional imperative content with satisfaction conditions. See Klein (2015) among others. I've criticized this view elsewhere (2018).

For convenience I will focus on pain experience in what follows.²

On Premise (1)

I will take the definition to be stipulative and well-motivated, theoretically as well as empirically. Even though it refers to *concepts* as mental representations in the psychologists' preferred sense, it can be seen as highlighting the essential role of successful or unsuccessful discriminatory response in the acquisition of any intentional content of sensory perception — whether that response is cognitive or behavioral. The full defense of premise (1) is beyond the scope of this paper, which requires going over certain areas of perceptual psychology and perceptual learning theory as well as theories of psychosemantics. But the basic idea is that, if it is not a mystery for sensory experiences to have representational content, what content an experience has or can have must be, to a large extent, a matter of what information can be extracted from the experience in the service of various cognitive and behavioral tasks downstream of the experience. Capacity for selective response to the informational content of an experience is therefore a necessary condition for that experience's having or acquiring a particular representational content. This capacity is a cognitive capacity in the broad sense, and I don't see any harm in characterizing it as a conceptual capacity. Without such a minimal requirement as expressed by (1) on the representation of *SP*, it is extremely difficult to see how reductive representationalists can hope to avoid making representational content of an experience a mystery.

There is also a general and quite an interesting issue about whether a more general version of premise (1) and the accompanying definition can plausibly be formulated to cover concepts that express less low-level properties that can legitimately enter into the content of experiences. By stipulation, I will take *SP* in (1) to be the lowest-level properties (roughly extensional with so-called traditional *secondary qualities* realistically interpreted) and remain agnostic whether *SP* can be extended to less low-level properties such as being a cube, or even a pine tree.

On Premise (2)

It is clear that the crux of defending this argument against reductive representationalists lies in defending premise (2). But luckily this premise is an empirical one, so settling it doesn't require a lot of fancy philosophical footwork. But it requires some clarification. What are the candidate proprietary concepts for pain experiences? What are the proprietary concepts in the case of color vision of a normally sighted person? These are more or less definite concepts that the person can non-inferentially apply to surfaces, volume, etc. while seeing

² But see my (2019a) for other experiences. The case of itches in some ways is even a stronger than pain. The concept of an itch on a bodily spot, when used correctly by the folk or scientists, doesn't attribute any physical condition of that bodily surface like some kind of disturbance. Hence ITCH, just like TICKLE and ORGASM, is not a proprietary concept with labeling uses — see below.

them, like the concepts RED, CRIMSON RED, DARK CRIMSON RED, *THAT* RED (as demonstrated), etc., attributing the relevant more or less specific sensible color properties.³ Similarly, the gustatory modality produces various more or less definite proprietary concepts that a taster can directly apply while tasting food items, like SWEET, SALTY, BITTER, etc. These would be used to attribute various sensible properties to these food items.

In the case of pain modality, nociception, the candidate proprietary concepts can be grouped under the quasi-technical term ‘painy’ as follows:

My toe is *painy* now IFF
a) my toe hurts, or
b) I feel/have (a) pain in my toe, or
c) there is (a) pain in my toe now.

I treat (a)–(c) to be roughly philosophically equivalent.⁴ A pain in my toe, then, is the instantiation of a property in my toe (or a “condition” of my toe) — the property of being painy. Hence, the candidate proprietary concept for nociception is PAINY if we want to make pain modality parallel to other modalities. Having clarified the second premise thus, we can still be somewhat relaxed about the proper expression of these concepts in language by denoting them by HURTS, PAINFUL, ACHES, PAINS, etc., depending on the context. The point is that despite the idiosyncrasies of natural languages, the canonical form of locating pain in body parts is purportedly a way of sensing the instantiation of a sensible property in that part, or some condition of that body part, and a first-person non-inferential application of the relevant proprietary concept on the basis of nociceptive experience is supposed to detect and identify that property or condition for downstream processing. We can now work on the special case of premise (2):

(2’) PAINY does not have any labeling uses.

This is a slightly fancy way of making the empirical claim that when people identify pains in their body parts, they don’t *ipso facto* identify a sensible property (or, condition) of those parts. If we interpret ‘sensible property’ as some objective (real) property or condition of body parts like some kind of tissue disturbance, (2’) is an empirical claim, and the evidence for it is decisive. The standard way to test claims like this is to collect evidence on the

³ I will put names of concepts in small cap. Thus, RED refers to the proprietary concept of red, and applies to red things (i.e., attributes the property of being red when activated).

⁴ I am pretty sure there are conversational contexts they aren’t, but here I follow Hyman’s (2003) hurt/pain paraphrasing procedure — more or less. Why not use ‘painful’ which, unlike ‘painy,’ is at least plain English? We could have used it, probably, if it hadn’t been adopted to mean something more specific in recent philosophy of pain, namely, to mean negative affect or hedonic tone that sensory pains typically come with. So, accordingly, ‘painful pain’ doesn’t express a redundancy.

basis of how people make pain attributions to body parts. To simplify things, consider John who, after coming from a long morning run, starts feeling a pain in his toe, and sincerely makes the following claim:

(L) I have a pain in my toe.

According to our framework, John has just expressed his thought (judgment) whose content is

(T) My toe is painful

If the concept PAINY had labeling uses, and thus were a proprietary concept for nociception, John's thought would be true or false according to whether or not John's toe has some kind of physical disturbance in it, assuming that some kind of disturbance is the sensible property or condition attributed by the concept. We could test this by investigating whether people would continue to assent to sincere utterances of (L) where we stipulate that there is no disturbance but they continue to feel pain in their toe as in well-defined so-called referred pain cases. Conversely, we could easily find out whether people would assent to (L) when there is the relevant kind of disturbance in their toe but they don't feel any pain there.

There is not a shred of credible empirical evidence that people use proprietary pain concepts with labeling uses. I won't belabor this point at length. I just want to point out that if there were using such concepts they would be treating a pain in a body part as one and the same as some physical disturbance in that part — they would be identifying pain and disturbance in the strict sense. As a matter of empirical fact, after reflection nobody does that. Not a single figure in the history of medicine, not a single philosopher in the history of philosophy has identified bodily pains with objective bodily disturbance.⁵ As far as I know, the only relationship acknowledged between them is some kind causation or correlation — pain as fairly direct sign or indication of some disturbance. According to early thinkers, this correlation was very strong, but it has become common sense in the last 50 years or so in pain science and clinical practice that the correlation is actually weak. In fact, it was the observation of this weak link that led to the big advances in pain science and clinical practice as anyone familiar with the history of pain science in the 20th century would readily acknowledge.⁶

⁵ With the exception of Alex Byrne (forthcoming). Rey's *The History of Pain* (1995) is a comprehensive and authoritative source for how pain conceptions and medical practices have evolved since ancient times. Dallenbach (1939) contains detailed discussion of the developments especially in the 19th and early 20th centuries.

⁶ Melzack and Wall's (1996) is still the best source of the changes in pain science that took place starting in the 1960's.

Adam Shriver and I have been collecting data about how pain scientists, clinicians, and other health professionals conceptualize attributions of pain to body parts in relation to physical stimuli. Since 2015 we have collected close to three hundred (anonymous) responses to one or another version of a questionnaire.⁷ We have found *no evidence whatsoever* that they think that pain in a body part is one and the same thing as some physical disturbance in that part. On the contrary, they seem explicitly, quite self-reflectively and vehemently opposed to making such an identification as has been emphasized again and again in the various versions of the IASP definition of pain since 1979.⁸

What is more interesting is that even among the recent philosophers of mind⁹ whose perceptualist or representationalist views *should* imply that pain in a body part *is* some objective condition of that part, we don't find this identification: they generally handle the location of bodily pain as an intentional location represented in the experience: pains are experiences that represent disturbances in bodily locations (bodily "sensations"), but the relevant *purportedly* proprietary concepts thereby deployed (e.g., PAINY, HURTS, PAIN THERE, etc.) have, mysteriously, no labeling uses. Hence these philosophers must be rejecting premise (1) without the slightest argument produced — I suspect, even without realizing that this is what it is they are doing.

Premise (2) of the Ur-Argument is empirical and the evidence for it is decisive.¹⁰ Still, a reductive representationalist may concede the point but suggest that we should be

⁷ The questionnaire can be found here: <<https://www.surveymonkey.ca/r/G8GHGD7>>

⁸ IASP (1979, 2020); Raja et al. (2020). See also Aydede (2017, 2019b) for extensive discussion of the history of the attempts to define and revise the IASP definition.

⁹ Armstrong (1968), Pitcher (1970), Tye (2006), Dretske (2006), Bain (2003, 2007), among others.

¹⁰ In this context I should say a few words on the experimental work done by Kevin Reuter and Justin Sytsma and his colleagues surveying the folk about pain attributions (e.g., Reuter et al. 2014, Sytsma & Reuter 2018, among other works). Like everybody else they observe that the folk conception of pain involves two threads that they call *bodily* conception and *experiential* conception. Reuter et al. claim that the results of their surveys show that the folk heavily prefer the bodily conception, (which they never explicitly articulate to the satisfaction of an onlooker who wants to know whether the folk strictly identify pains in body parts with some physical conditions of those parts). I am not at all convinced that they show anything of the sort. Everybody agrees that the folk conception of pain has these two apparently conflicting (bodily and experiential) strands in it (see my 2005/2019 for explicit statement and elaboration) and that their open co-existence is *prima facie* puzzling. Depending on how you set up the vignettes and questions, you can get results supporting either strand — this is not surprising at all. To see a detailed and pretty convincing criticism of their work along these lines, see Borg et al. (2019). But unlike most everybody else, Reuter et al. think that these two conceptions are incompatible with each other in the sense that they can't be true at the same time. Indeed, following Chris Hill (2006), they present the tension as a *paradox*. This is extremely unfortunate since it involves attribution of a colossal inconsistency not only to the folk, but to everybody including pain scientists, clinicians, and other health professionals — not to mention scores of historical as well as contemporary philosophers and thinkers. This might have been OK if the inconsistency were deep and hidden somewhere. But these two conceptions have been in the open as long as people have been talking about pain. For

revisionist about our thinking and pain language: we *should* strictly identify a pain in a body part with some physical disturbance in that part, and *should* think and talk accordingly. Given the way basic and clinical pain science has been developing in the last 60 years or so, this will *not* happen. Furthermore, there is strong reason to believe that it *cannot* happen, empirically. The revisionist claim comes close to claiming that we should change the way our cognitive architecture works. If the way we, as a matter of fact, in the streets, labs or clinics, naturally cognitively resist identifying pain in a body part with an objective condition of that part, we should ask why. Why is it that, if nociception were strongly representational, we find an empirical pattern that is precisely the opposite of what is predicted? If nociceptive experiences had strong representational content, we would expect this to be reflected in the cognitive architecture that interfaces the incoming sensory information with early conceptual systems that works just like in other sensory modalities. But that is not what we find. Our cognitive architecture is obviously wired differently in nociception in not giving rise to proprietary concepts with labeling uses. It's not quite clear what the wisdom of suggesting a revisionist approach would be except saving a philosophical dogma. Besides, it isn't like we would be missing anything if we didn't identify pains with disturbances: we already have all the necessary conceptual and linguistic apparatus to think and talk about *disturbances* and the like — in the streets, in research labs and clinics.

If pain experiences were completely transparent to their owners about what they represent, we would have proprietary concepts of what they represent. But we don't have such concepts. Hence pain experiences are not strongly representational. This is not to say that pain and other bodily experiences are completely devoid of any intentional features. Clearly these experiences contain information about where in the body these sensations (*sic!*) occur — maybe along with their temporal features. If bodily location should be treated *intentionally* (as somewhat *denoted* or referenced by the experience), then the main question is about whether the same experience *attributes* (or predicates) any sensible properties or qualities to that location. The main contention of this paper is that for there to be such an attribution — for there to be representational content with accuracy conditions — proprietary concepts are required. This is premise (1). Premise (2) is the empirical claim that we don't have such concepts.

instance, the *Note* to the IASP definition of pain had contained this “clarification” until the last revision in 2020: ‘[pain] is unquestionably a sensation in a part or parts of the body, but it is also always unpleasant and therefore also an emotional experience’ (IASP 1979). But under the assumption of incompatibility, any evidence for bodily conception becomes evidence against the experiential conception, and vice versa. But if these conceptions are not incompatible — as should have been the obvious default assumption — then finding evidence for the bodily conception or against the experiential conception doesn't cut any ice. If these conceptions are not incompatible, then this means that there is philosophical work to be done to understand how these two apparently conflicting conceptions manage to peacefully share the same bed so intimately. The empirical surveys of folk intuitions will not settle the matter. The experimental work of Reuter et al. is deeply problematic, so I set it aside.

One might inquire that in our conceptual response if we don't locate physical disturbances in having pain experiences, what do we locate, if anything? This is a good question, and the answer is not obvious.¹¹ One might think that irrealist non-reductive versions of strong representationalism may have a better chance of answering this question. In the remainder of this paper, I will, rather more briefly, look at whether this is so.

Irrealist Representationalism

The version I'll develop will be schematic as there are various versions of the view, and none of them seems to have gained the status of the canonical formulation.¹² According to irrealist representationalism, the "sensible" *qualities* (non-conceptually) represented in sensory experiences are nowhere instantiated in the world in which the experiencer lives — not in the world and not in the experience. These qualities live only in the intentional content of the experiences. Sensory experiences represent these qualities as instantiated in the world, of course. So these experiences are always non-veridical. If my visual experience represents this ripe tomato as pure red, even under perfectly normal conditions, my visual experience is in error, the tomato is anything but pure red. Pure red, as given to me in my experience, is simply the kind of quality that cannot have instances in this world. Even though no experiences of color-sighted people can ever be veridical (in this world), color experiences still *represent* these *sui generis* primitive qualities as instantiated. This view is sometimes called *figurative projectivism*,¹³ and like all projectivist views, it is an error-theory. The sensible world as it is in itself is nothing like as it appears to me in my sensory experience, but because the error is systematic, it matters little for the success of my behavior as I negotiate my way in the world.

Defenders of irrealist strong representationalism tend to be non-reductivist about both experience and representation: sensory experiences have their phenomenal character solely due to their representational content (sometimes called *phenomenal content*), which cannot be explained in purely non-representational, non-mental, or more broadly, naturalistic/physicalistic terms. Partly because of this, it is harder to criticize them by pointing out that they are making a mystery out of an experience's representational powers. Nevertheless, representationalism about sensory experiences doesn't come cheap, and we can run a version of the Ur-Argument against them.

(1*) For every sensory modality (or submodality) *m* that gives rise to conscious experience *e* in that modality, there is a set of low-level sensory qualities *SQ* associated with *m* such that *e* represents *SQ* only if *e* is the immediate epistemic basis for the acquisition and deployment of the *proprietary concepts* of *SQ*.

¹¹ But see my (2019a and 2020) for a more comprehensive discussion and an answer.

¹² Here I have in mind views of the kind Chalmers (2006) and Pautz (2010) defend — there are of course others.

¹³ See Shoemaker (1994).

(Definition*)

c is a *proprietary concept* of *SQ* IFF

- (i) *c* is needed for the explanation of how the agent *g* (where *c* belongs to *g*'s set of mental representations) discriminates or learns to discriminate instances or quasi-instances of *SQ*, and
- (ii) *c* is necessarily acquired from sensory experiences *e* in *m*, and
- (iii) *c* has **labeling** uses in that they are non-inferentially (directly/immediately) applied by *g* to the instances or quasi-instances of *SQ* on the basis of *e* — thus generating true or false *de re* thoughts non-inferentially based on veridical or non-veridical sensory episodes.

The rest of the argument is the same except the conclusion (5) is now about the strong representationalism of the irrealist variety. I will call the resulting argument against irrealist representationalism as such, the **Extended Ur-Argument**.

A *quasi-instance* of a sensory quality *SQ* represented in sensory experiences is the instance of a sensible *property SP* that is the normal cause of experiences that (non-conceptually) represent *SQ* (as instantiated).

This needs a bit of unpacking. Suppose you are an irrealist representationalist. You believe that the color of the ripe tomato in front of you that you see is not in fact instantiated on the surface of the tomato. This is the sensory *quality* as given to you in your experience — it's an *Edenic quality*, in Chalmers' terminology (2006). Although it appears instantiated by the tomato, the appearance is non-veridical. What is instantiated in fact by the tomato is a physical sensible *property (SP)*, something like a surface spectral reflectance (*ssr_{red}*) of a certain kind, that is the normal cause of experiences that represent (as instantiated) the sensory *quality* — the Edenic quality (*SQ*). This quality, although represented as instantiated in the world, is not instantiated anywhere, not even in the sensory episode that you undergo that *is* the event of your sensorially representing the quality. So, these Edenic sensory qualities are not *sensible* qualities. There are no instances of them — they *cannot* be instantiated — in the world you live. So, they cannot be sensed — they can only be represented as instantiated non-veridically. Hence the new terminology. We will continue to refer to the real objective sensible properties instantiated in the world that are the normal causes of relevant sensory experiences as such, as *sensible properties*. We will use the term *sensory quality* for those Edenic qualities directly represented in the experience.

Now depending on the theorist, there will be different kinds of complicated relationships among the sensory experiences, the thoughts that these experiences will non-inferentially prompt, and the expressions of these thoughts in the natural language of the experiencer.

Nevertheless, these theorists cannot afford to ignore empirical facts about how we in fact think and use the language.

So you are looking at the ripe tomato, and the following thought occurs to you:

<Ah, that's a lovely red!>

where 'that' refers to the surface of the tomato. Let's stipulate that you have the following thought tokened in your mind immediately as a result of your visual experience of the tomato:

#RED(*that*)

This is a convenient notation depicting the format of your thought and its elements. '#' indicates active deployment (in this case, labeling use). If you are an irrealist representationalist, you have a few options about what to say regarding the truth-value of this thought.

- A. #RED(*that*) attributes the Edenic sensory quality to the surface of the tomato and is therefore false.
- B. #RED(*that*) attributes the sensible property *ssr_{red}* and is therefore true (on this occasion).
- C. #RED(*that*) has two sets of truth-conditions associated with it: strict and loose.¹⁴ Strictly, it is false, but loosely it is true. Strict truth-conditions are usually in play in philosophical contexts. It is the loose truth-conditions that apply in pretty much every other contexts.
- D. #RED(*that*) has a layered neo-Fregean content. Its primary intension is something like: <Primitive redness is instantiated there>. This thought is always false in the actual world. But given the functional role of this color experience (something like <the normal cause of this experience representing the instantiation of primitive red there is instantiated there>), this may have different extensions in different worlds. When evaluated in our world, this thought is true due to the tomato's *ssr_{red}* causing your experience on this occasion (*ssr_{red}* being the normal cause red representing experiences in the actual world).

C and D can be related by saying that D provides a more philosophically satisfying interpretation of C. So much for the schematic presentation of irrealist representationalism. There is a lot more to say, especially about how such thoughts are supposed to connect up

¹⁴ Somewhat similar to the perfect and imperfect truth-conditions Chalmers (2006) introduces for experiences.

with the standard truth-values of utterances expressing these thoughts in public languages. But the above will be sufficient to run the extended Ur-Argument against the view.

So what happens to John's thought prompted by his feeling pain in his toe?

(T) #PAINY(*that*)

where '*that*', in the language of his thinking, refers to his toe where John feels the pain.

Option A, when transferred to the case of pain, results in the wrong verdict: (T) is intuitively true. John may not have much idea about the truth-makers of his thought. But just as John himself would insist, I am pretty sure that whatever this thought attributes to the toe, John's thought is true.¹⁵ Any view that implies that John's thought is in fact false and that therefore John doesn't have a pain in his toe is, given empirical facts, therefore refutes itself.

Option B also results in the wrong verdict: as argued before, when we think of a pain in a body part, we don't *ipso facto* think any objective condition of that body part like a disturbance. #PAINY doesn't attribute a *sensible property*. In fact, let us stipulate that John's toe has no disturbance of the relevant sort in his toe at all. The disturbance, let's say, is in his lower spine. John's pain in his toe is now a referred pain. So his thought is still true: John correctly judges that he has a pain in his toe.

If options A and B result in the wrong verdict, then combining them will obviously not help. So option C is no good either. To see this, consider John's case when he sincerely judges that he has a pain in his toe, even when there is no disturbance of the relevant kind in his toe. The empirical facts are such that people (pretty much everyone from common folk to scientists/clinicians, and even philosophers) consider his belief to be true. But option C gives the wrong verdict by requiring that John's belief is false. The response to option D follows the same path: in a case where the pain is referred, for instance, John's thought turns out to be false.

Now I'm sure, given their metaphysical resourcefulness, irrealist non-reductive representationalists will find a way to resist this argument. Perhaps they will insist that option A gives us the right verdict. The Edenic pained quality is nowhere instantiated, let alone in body parts. Still it's falsely represented by pain experiences as so instantiated. So

¹⁵ Maybe, #PAINY *correctly* attributes an Edenic quality to the toe after all. If this turns out to be so, then the metaphysics of pains and other similar bodily sensations is constituted by the instantiation of these qualities in the actual world. So maybe some such represented sensory qualities are instantiated after all. In fact there is a strong pull in the common sense conception of pain in this direction. Not to mention some philosophers too: Jackson (1977, among other sense-datum theorists), Hyman (2003), Bradley (2021), among others. However, such a view would be intolerably mysterious. My own views don't appeal to any primitive or Edenic qualities — see my (2005/19, 2019b, 2020).

whatever John may sincerely think (deploying PAINY) directly prompted by his experience, he doesn't in fact have a pain in this toe (no sensory quality is instantiated in his toe): he is simply deluded in believing that he has a pain in his toe. Note that this is quite different from the case where John judges, directly prompted by his seeing red and deploying RED, that this (tomato) is red. Here there is room for the irrealist to claim that John's thought is strictly speaking false (the Edenic red is nowhere instantiated), but is loosely true due to the instantiation of *ssr_{red}* on the tomato's surface. If it turns out that the tomato is not in fact red (no *ssr_{red}*), then it is not unreasonable to think that John's thought is in fact false even loosely speaking. But no such room exists in the case of John's belief that he has a pain in his toe.¹⁶

At this point, I can't do any better than giving the representationalists an incredulous look if they indeed insist that A gives us the right verdict. Having become pretty immune to incredulous looks, I don't, however, expect irrealist representationalists to be moved by that.

The best move on their part, it seems to me, is to reject premise (1) of the extended Ur-Argument on the ground that some sensory experiences can represent certain sensory qualities even though this representation never gives rise to labeling concepts, thus never generating relevant range of thoughts with truth-conditions of any sort. We have concepts of these, yes, but they only allow us to think about these qualities, in the abstract so to speak, without these concepts ever getting applied to particulars — appearances notwithstanding. These qualities are never instantiated anywhere in the world of the thinker after all. So these concepts in fact never get applied to particulars represented by the relevant sensory experiences. But this seems just empirically false. It would make a bad joke in the light of billions of people experiencing bodily pains in their daily routine — not to mention millions of chronic pain sufferers — correctly thinking and reporting pains in their body parts. How could they be thinking about bodily pains in the abstract, without judging they have pains in body parts? It is one thing to claim that these concepts are applied (have labeling uses) but they always generate false thoughts (as in option A); it's another to claim that they are, as a matter of fact, never put in (apparently) labeling uses — they are there to merely express what the content of their experience is. Besides, the very idea of an experience as a representation demands that it's the sort of thing that can be put in psychological, epistemological, behavioral use *about what they represent*. How is that possible without these experiences giving rise to proprietary concepts?

A representationalist at this juncture is likely to realize that what they want to claim about the truth-makers of (T) is not that <John's toe instantiates a quality/property>, but that <John's experience represents the instantiation of a certain quality/property in his toe>. But note that if the concept PAINY attributes the property of *being an experience representing a*

¹⁶ This is not because I believe in the incorrigibility of these beliefs and their reports. Rather, it has to do with their truth-conditions and what the empirical facts are.

sensory quality or sensible property, it is not a proprietary concept for the pain modality. It's in fact an intentional property not attributed to the toe at all. It is the old concept of a *pain experience* implicating a toe (a *pain sensation* in the toe). When John judges (T), he is making a self-attribution: <I'm experiencing this quality as if it's in my toe>. On this proposal, note that 'this quality' does not refer to the pain in his toe — contrary to what John himself believes. Rather it's supposed to refer to the *apparent* instantiation of the primitive sensory quality. But the reference fails, of course: it's a use of an empty mental demonstrative — there are no instances to be demonstrated according to irrealists. We are not only missing proprietary concepts for pain experiences to be representational as per premise (1), but under this proposal, we are also burdened by the fact that the sensory quality John is claimed to be presented with in his experience is not in fact the pain quality he seems to be locating in his toe. So, flat-out denying this premise won't solve the mystery of missing concepts for the irrealist: if PAINY attributes an intentional property, we don't seem to have a concept by which to express what sensory quality is represented in pain experiences — even if this concept is not a proprietary one. And worse, very counterintuitively, this represented quality turns out to be *not* the pain quality itself.¹⁷

I conclude that making pain experiences (along with the experiences of other similar bodily sensations) strongly representational is a mistake. These *experiences* don't attribute any property or quality to body parts. But if pains are not strongly representational, then there are non-representational sensory experiences. But strong representationalism claims that all phenomenologically salient sensory experiences are strongly representational. We need to find new alternatives.¹⁸

References

Armstrong, David M. 1968. *A Materialist Theory of the Mind*. New York: Humanities Press.

Aydede, Murat. 2005/2019. Pain. *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/spr2013/entries/pain/>>

¹⁷ Chalmers (2006) identifies the property of *sensorially representing Edenic red* (in having a visual experience) with a phenomenal property, *phenomenal red*. In this terminology, John's concept PAINY would be attributing a phenomenal property (thus, the concept would be a phenomenal concept), and (T) would express the self-attribution of this property. In this case, PAINY would not be proprietary, of course. What is puzzling in this case, though, is that John (along with others) would be missing a concept with which to express the Edenic property that the phenomenal property is supposed to represent. One would have thought that the pain sensation/quality that appears to be instantiated in the toe is this Edenic quality represented in the experience (partly constituting the phenomenal property), but, on this proposal, John's thought doesn't attribute this property to his toe as expected.

¹⁸ I started to develop an alternative to strong representationalism in Aydede & Güzeldere (2005), and my (2019a, 2020). I would like to thank Dom Alford-Duguid, Jonathan Cohen, Brian Cutter, and Matt Fulkerson for very useful conversations and comments on earlier versions.

- Aydede, Murat. 2017. Defending the IASP Definition of Pain. *The Monist* 100 (4): 439–64.
- Aydede, Murat. 2018. Critical Notice of Colin Klein's *What The Body Commands: The Imperative Theory of Pain*. "Author meets Critics" session in the Eastern APA 2015, Baltimore, MD. (Available online: <<https://philpapers.org/rec/MURCNO>>)
- Aydede, Murat. 2019a. Is the Experience of Pain Transparent? Introspecting Phenomenal Qualities. *Synthese* 196: 677–708.
- Aydede, Murat. 2019b. Does the IASP definition of pain need updating? *PAIN Reports*, 4(5), 1–7.
- Aydede, Murat. 2020. What Is a Pain in a Body Part? *Canadian Journal of Philosophy*, 50(2), 143–158.
- Aydede, Murat, and Güven Güzeldere. 2005. Cognitive Architecture, Concepts, and Introspection: an Information-Theoretic Solution to the Problem of Phenomenal Consciousness. *Noûs* 39 (2): 197–255.
- Bain, David. 2003. Intentionalism and Pain. *Philosophical Quarterly* 53(213): 502–522.
- Bain, David. 2007. The Location of Pains. *Philosophical Papers* 36(2): 171–205.
- Borg, Emma, Richard Harrison, James Stazicker, and Tim Salomons. 2019. Is the folk concept of pain polyeidic? *Mind and Language*, online first. DOI: 10.1111/mila.12227.
- Bradley, Adam. 2021. The Paradox of Pain. *The Philosophical Quarterly* 71 (4): 717–37.
- Byrne, Alex (forthcoming). Knowing that I am in pain. In *Self-Knowledge and Knowledge A Priori*, edited by Giacomo Melis, Giovanni Merlo, Crispin Wright, Oxford UP.
- Chalmers, David J. 2006. Perception and the fall from Eden. In *Perceptual Experience*, edited by Tamar S. Gendler & John Hawthorne, 49–125. Oxford: Oxford University Press.
- Dallenbach, K. M. 1939. Pain: history and present status. *The American Journal of Psychology*, 52(3), 331–347.
- Dretske, F. 2006. The Epistemology of Pain. In M. Aydede (Ed.), *Pain*. (pp. 59–73). MIT Press.
- Hill, Christopher. 2006. Ow! The Paradox of Pain. In *Pain: New Essays on Its Nature and the Methodology of Its Study*, edited by Murat Aydede, 75–98. Cambridge, Massachusetts: MIT Press
- Hyman, John. 2003. Pains and Places. *Philosophy* 78 (1): 5–24.

IASP (International Association for the Study of Pain). 1979. IASP Pain Terms — 1979. *Pain*, 6: 247–252. (Revised on-line version: <<http://www.iasp-pain.org/Education/Content.aspx?ItemNumber=1698>>, 2020)

Jackson, Frank. 1977. *Perception*. Cambridge UP.

Klein, Colin. 2015. *What the Body Commands: The Imperative Theory of Pain*. MIT Press.

Melzack, Ronald, and Patrick D. Wall. 1996. *The Challenge of Pain*. Updated second edition. Penguin.

Pautz, Adam. 2010. Do Theories of Consciousness Rest on a Mistake? *Philosophical Issues* 20: 333–67.

Pitcher, George. 1970. Pain Perception. *The Philosophical Review* 79(3): 368–393.

Raja, S. N., Carr, D. B., Cohen, M., Finnerup, N. B., Flor, H., Gibson, S., et al. (2020). The revised International Association for the Study of Pain definition of pain: concepts, challenges, and compromises. *Pain*, 161(9), 1976–1982.

Reuter, K, D. Phillips, and J. Sytsma. 2014. Hallucinating Pain. In *Advances in Experimental Philosophy of Mind*, edited by Justin Sytsma. London: Bloomsbury.

Rey, R. 1995. *The History of Pain*. Cambridge, MA: Harvard University Press.

Shoemaker, Sydney. 1994. Phenomenal Character. *Noûs* 28 (1): 21–38.

Sytsma, J., and K. Reuter, K. 2018. Unfelt pains. *Synthese*. Online First.

Tye, Michael. 2006. Another Look at Representationalism about Pain. In *Pain: New Essays on Its Nature and the Methodology of Its Study*, edited by Murat Aydede, (pp. 99–120). Cambridge, Massachusetts: MIT Press.